

## **Ex. 3-B**

# **Special Master's Amended Order Regarding Production of Named Plaintiff Data and the Exhibits thereto**

## **Redacted Version of Document Sought to be Sealed**

**UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA  
SAN FRANCISCO DIVISION**

IN RE: FACEBOOK, INC. CONSUMER  
PRIVACY USER PROFILE LITIGATION,

This document relates to:

ALL ACTIONS

**MDL NO. 2843**

CASE NO. 3:18-MD-02843-VC-JSC

HON. VINCE CHHABRIA  
HON. JACQUELINE SCOTT CORLEY  
COURTROOM 4 – 17<sup>TH</sup> FLOOR  
SPECIAL MASTER, DANIEL GARRIE, ESQ.

**AMENDED ORDER REGARDING  
PRODUCTION OF NAMED PLAINTIFF  
DATA**

**INTRODUCTION**

1. Pending before the Special Master is Plaintiffs' Motion to Compel Production of Named Plaintiffs' Content and Information.

**BACKGROUND**

2. On November 25, 2019, Plaintiffs' served Requests for Production Nos. 9-13, which seek documents relating to the named Plaintiffs in this matter ("Named Plaintiffs"). See Exhibit A (RFPs 9-13). In brief, Request No. 9 seeks all documents relating to each of the Named Plaintiffs; Request No. 10 seeks documents sufficient to show the categories of content and information Facebook collects, tracks, and maintains about them; and Requests Nos. 11-13 seek documents identifying third parties that were able to access information about the Named Plaintiffs. Id.

3. In response to Requests for Production Nos. 9-13, Facebook produced more than one million pages of individual user data it maintained relating to the Named Plaintiffs, most of which was obtained from the "Download Your Information" tool ("DYI Tool"). The data obtained from the DYI Tool is mostly limited to information pertaining to users' on platform Facebook activity. See Exhibit B (DYI Data).

4. Statements by Facebook's counsel during an August 14, 2020, discovery hearing indicated that Facebook maintained additional data related to the Named Plaintiffs that was not produced. See Exhibit C (8/14/2020 Discovery Hearing Transcript) at 8:10-13 ("There is other – there's Facebook-generated information, information generated by third parties, information received from third parties. We have not represented that that is comprehensively included in our production.").

5. Plaintiffs filed a motion in September 2020 to compel additional discovery related to Requests for Production Nos. 9-13. See Exhibit D (Motion to Compel Additional Discovery Related to RFPs 9-13). Plaintiffs asked the Court to compel production of sensitive information Facebook derives and collects from business partners, app developers, apps, and other sources. This request included

1 "native, appended and behavioral data" and purportedly anonymized data that could be connected to the  
2 Named Plaintiffs. Id. at 7-11.

3 6. On October 8, 2020, Facebook responded to Plaintiffs' motion to compel. See Exhibit E  
4 (Facebook's October 2020 Response). Facebook contended that all information related to the Named  
5 Plaintiffs that they did not themselves share on Facebook was outside the scope of the case; that all  
6 information not shared through one of the four theories of the case was not within the scope of the case;  
7 that Plaintiffs were not entitled to all data collected from third parties about the Named Plaintiffs; that  
8 the Stored Communications Act and Video Protection Privacy Act claims did not require the production  
9 of additional data Facebook had collected about the Named Plaintiffs; and that Facebook could not  
10 reasonably collect any of the additional information Plaintiffs sought. Id. at 6-10.

12 7. On October 29, 2020, Judge Corley issued Discovery Order No. 9, ruling "that discovery is  
13 not as limited as Facebook contends" and "the discoverable user data at issue includes: [1] Data  
14 collected from a user's on-platform activity; [2] Data obtained from third parties regarding a user's off-  
15 platform activities; and [3] Data inferred from a user's on or off-platform activity." See Exhibit F  
16 (Discovery Order No. 9) at 2.

18 8. Facebook did not produce additional documents in response to Requests for Production Nos.  
19 9-13.

20 9. On October 6, 2021, Special Master Garrie and Judge Andler declared impasse on the issue  
21 of whether Facebook should be compelled to produce additional documents related to the Named  
22 Plaintiffs pursuant to Discovery Order No. 9.

24 10. On November 29, 2021, Special Master Garrie issued an Order Re: Plaintiffs' Motion to  
25 Compel Production of Plaintiff Data which found that "Discovery Order No. 9 does not limit the scope  
26 of discoverable data related to the Named Plaintiffs to data that was shared with third parties, as  
27 Facebook contends, because Judge Corley's ruling contains no language indicating such a limitation."  
28 See Exhibit G (Order Re: Plaintiffs' Motion to Compel Production of Plaintiff Data) at 4. The November

1 29, 2021, order also required Facebook to provide a list of data sources that may contain Named  
2 Plaintiff data, including descriptions of the data sources and the Named Plaintiff data they may contain.

3 11. Facebook subsequently provided a list of data systems that may contain Named Plaintiff data  
4 without descriptions of the systems or data. See Exhibit H (Declaration of David Pope, Exhibit A).

5 12. On December 29, 2021, Special Master Garrie issued an Amended Order Re: Plaintiffs'  
6 Motion to Compel Production of Plaintiff Data, which required Facebook to provide descriptions of the  
7 purposes of each system and the business units that use each system. Special Master Garrie held a  
8 hearing with David Pope to address these points. See Exhibit I (David Pope Hearing Transcript).

9 13. Over the next three months, Special Master Garrie held a series of hearings with various  
10 Facebook engineers and requested documentation to develop an understanding of the systems identified  
11 by David Pope. The submissions and findings in connection with these hearings are reflected in the  
12 following documents: Exhibit J (Facebook's January 6, 2022 Letter); Exhibit K (Facebook's January 27,  
13 2022 Letter); Exhibit L (February 17, 2022 Hearing Transcript); Exhibit M (Facebook's March 7, 2022  
14 Letter); Exhibit N (March 9, 2022 Hearing Transcript).

15 14. On March 22, 2022, Special Master Garrie issued the Order Following March 9, 2022,  
16 Hearing Regarding Plaintiffs' Motion to Compel Production of Plaintiff Data, which required each party  
17 to submit a proposed protocol for production of Named Plaintiff data. See Exhibit O (March 9, 2022,  
18 Hearing Order).

19 15. The parties subsequently submitted their proposed protocols and on May 17, 2022, Special  
20 Master Garrie held a hearing with the parties to identify and resolve areas of disagreement with respect  
21 to proposed protocols. See Exhibit P (Facebook's Letter of April 18, 2022) and Exhibit Q (Plaintiffs'  
22 Letter of April 29, 2022).

23 16. Following the May 17, 2022 hearing, the parties agreed on all aspects of the proposed  
24  
25  
26  
27  
28

1 protocols except for the following three issues: selection of Hive tables for production;<sup>1</sup> production of  
2 data related to apps installed by friends of Named Plaintiffs; and whether Facebook should search cold  
3 storage for Named Plaintiff data in Hive. Special Master Garrie requested briefing on these three issues,  
4 an updated proposed protocol submission from Facebook reflecting the areas of agreement, and  
5 additional information from Facebook regarding the Hive tables. See Exhibit R (May 17, 2022 Hearing  
6 Order); Exhibit S (Facebook's May 30, 2022 Letter); Exhibit T (Facebook's June 2, 2022 Letter).

7  
8 17. On June 7, 2022, the parties submitted their briefs on the three outstanding issues regarding  
9 the Named Plaintiff data proposals.

10 18. Facebook proposed searching for and producing Named Plaintiff data from a sample of the  
11 11,051 Hive tables identified in Exhibit B to Facebook's April 11, 2022 submission.<sup>2</sup> See Exhibit U  
12 (Facebook's June 7, 2022 Letter). The sample would include 250 tables selected by Facebook and 250  
13 tables selected by Plaintiffs. Id. To assist Plaintiffs in the table selection, Facebook agreed to provide  
14 Plaintiffs with the schema (column names) for all 11,051 Hive tables from which the sample would be  
15 selected. See Exhibit T (Facebook's June 2, 2022 Letter). Facebook agreed to "provide Plaintiffs data  
16 regarding interactions that friends of the named plaintiffs had with businesses/apps using Facebook  
17 Login, without identifying which friend interacted with each business/app." See Exhibit U (Facebook's  
18 June 7, 2022 Letter). With respect to the cold storage issue Facebook stated that "Facebook is willing to  
19 consider restoring data from cold storage, but identifying what, if any, data should be restored is  
20 premature at this stage." Id. Facebook instead proposed the following: "Once the Hive tables have been  
21 identified, Facebook will evaluate which tables, if any, include data in cold storage, and make a proposal  
22 regarding what, if any, data it can reasonably restore, search, and produce." Id.  
23  
24

25  
26 <sup>1</sup> Facebook initially proposed searching a sample of 200 Hive tables (100 selected by Plaintiffs and 100 selected by  
27 Facebook) for Named Plaintiff data. See Exhibit P (Facebook's Letter of April 18, 2022). Plaintiffs initially proposed that  
28 Facebook produce the first five rows of all 11,051 Hive tables identified in Exhibit B to Facebook's April 11, 2022  
submission. See Exhibit Q (Plaintiffs' Letter of April 29, 2022).

<sup>2</sup> Facebook identified these 11,051 Hive tables as potentially containing Named Plaintiff data using a data classification tool  
that can classify tables within Hive as containing UID, RID, or SID (identification numbers for individual Facebook users).  
See Exhibit T (Facebook's June 2, 2022 Letter).

1 19. Plaintiffs' June 7, 2022 submission largely did not address the three issues on which Special  
2 Master Garrie requested briefing in the May 17, 2022 Hearing Order. Instead, Plaintiffs argued "new  
3 evidence has come to light in two 30(b)(6) depositions related to those questions" showing that (1)  
4 Facebook selected 137 Hive tables and put them in "cold storage" precisely because they were relevant  
5 to this litigation; (2) Facebook is capable of searching offline Hive tables using HQL and the Scuba tool;  
6 (3) the DYI file is not the most complete or usable compilation of user data; and (4) Facebook has  
7 withheld from production at least 52 snapshots of Named Plaintiff data using a never-before revealed  
8 tool more commonly used to collect user data called Switchboard. See Exhibit V (Plaintiffs' June 7,  
9 2022 Submission).

11 20. At Special Master Garrie's request, the parties submitted additional briefs on the issues  
12 regarding the 137 Hive tables Facebook preserved in cold storage and the Switchboard snapshots. See  
13 Exhibit W (Facebook's June 16, 2022 Letter); Exhibit X (Plaintiffs' June 20, 2022 Letter).

14 21. Facebook responded to Plaintiffs' claims regarding preservation of the Hive tables as  
15 follows:

17 Tables in Hive were preserved in connection with this litigation and related litigation for  
18 a number of reasons, many of which have nothing to do with Named Plaintiff data. Of the  
19 137 tables identified, a minority contain user identifiers, and several of those tables were  
20 put on hold in connection with Facebook's April 18, 2022 proposal to produce categories  
of data requested or referenced by Plaintiffs. See Exhibit W (Facebook's June 16, 2022  
Letter) at 3.

21 Facebook further stated that it agreed to produce the schema (i.e. column names) for the preserved tables  
22 containing user identifiers and will meet and confer with Plaintiffs regarding a production of Named  
23 Plaintiff data from those tables. Id. at 3. Facebook also agreed to produce Named Plaintiff data from the  
24 Switchboard system. Id. at 2. Facebook stated that Facebook had not previously identified or produced  
25 data from Switchboard because the data in Switchboard is largely a subset of the data in the DYI files  
26 already produced. Facebook stated that the data contained in Switchboard and not DYI includes "(1)  
27  
28

1 data about other users and (2) data created for law enforcement (neither of which are relevant to these  
2 Named Plaintiff data proceedings)." Id. at 1-2.

3 22. Plaintiffs argue, among other things, that representations made by Facebook's counsel with  
4 respect to Switchboard and the Hive tables is inconsistent with Facebook's sworn testimony. See Exhibit  
5 X (Facebook's June 20, 2022 Letter) at 2. Plaintiffs request an order requiring a production date of  
6 Friday, June 24, 2022 for the Hive schema, fields and documents sufficient to describe the contents of  
7 the tables. Id. at 2. Plaintiffs further request an order to show cause why Facebook should not be ordered  
8 to produce all Hive tables referencing user identifiers immediately. Id. at 2. Plaintiffs also argue that  
9 internal Facebook communications produced to date indicate that it is easier for Facebook to restore data  
10 from cold storage than Facebook has represented, and Facebook should be ordered to restore data from  
11 cold storage. Id. at 3.

### 13 FINDINGS

14 23. Special Master Garrie finds that the parties agree with respect to the production of the  
15 following buckets of data: (1) user objects and associations to those objects in the TAO system for each  
16 Named Plaintiff; (2) Named Plaintiff data in the Switchboard system; (3) specific types of data in Hive  
17 requested or referenced by Plaintiffs in challenging Facebook's production of Named Plaintiff data;<sup>3</sup> and  
18 (4) data regarding interactions that friends of the Named Plaintiffs had with businesses/apps using  
19 Facebook Login (without identifying which friend interacted with each business/app). See Exhibit S  
20 (Facebook's May 30, 2022 Letter); Exhibit U (Facebook's June 7, 2022 Letter).  
21

22 24. Special Master Garrie finds that the parties agree on the following additional items in  
23 connection with the production of Named Plaintiff data: (1) Facebook will produce the TAO schema for  
24 the TAO data to be produced; (2) Facebook confirms it will produce the Hive data described in  
25 paragraph 23 above regardless of whether it appears in the DYI files; (3) Facebook will provide the  
26  
27

---

28 <sup>3</sup> This includes off-platform activity, ad interests, ad click data, ad impressions data, and custom audience data. See Exhibit S (Facebook's May 30, 2022 Letter) at 3.



1 names of the tables from which the Hive data described in paragraph 23 above will be produced, how  
2 Facebook identified the tables, and the schema for such data; (4) Facebook will provide the schema  
3 (column names) for all 11,051 Hive tables identified in Exhibit B to Facebook's April 11, 2022  
4 submission and the schema for all of the 137 preserved Hive tables that contain user identifiers; (5)  
5 Facebook will produce "set permissions" (audience controls on a post), including any audience controls  
6 on individual pieces of content, from TAO; (6) Facebook will produce updated privacy settings for each  
7 Named Plaintiff. See Exhibit S (Facebook's May 30, 2022 Letter).

8  
9 25. The issues on which the parties do not appear to agree are the procedure for sampling Hive  
10 tables and searching cold storage for Hive tables.

11 26. With regards to Hive table sampling, Facebook's proposal to provide the schema for the  
12 11,051 Hive tables, search a sample 500 tables (250 selected by Facebook and 250 selected by  
13 Plaintiffs) for Named Plaintiff data, and produce Named Plaintiff data identified is more appropriate  
14 than Plaintiffs' initial proposal of requiring Facebook to produce the first five rows of all 11,051 Hive  
15 tables because the Hive table schema Facebook agreed to provide would give Plaintiffs substantially the  
16 same information sought by Plaintiffs' proposal without the additional burden of producing tens of  
17 thousands of rows of data that are unlikely to contain Named Plaintiff data. Plaintiffs also will be able to  
18 use the Hive table schema to select tables they believe are most relevant.

19  
20 27. With regards to searching cold storage for Hive tables, Facebook's proposal to first identify  
21 the Hive tables to be searched, then determine whether any tables contain data in cold storage, and, if so,  
22 make proposals as to searching cold storage is appropriate, as we will be better positioned to develop a  
23 protocol for searching cold storage once we know the volume of cold storage data there will be for the  
24 Hive tables selected (if any). Facebook's proposal also preserves both parties' rights to make arguments  
25 regarding cold storage once the tables are identified.

26  
27 28. Based on Special Master Garrie's interviews with Facebook engineers and the Facebook's  
28

1 submissions regarding the other systems identified in the Exhibit H (aside from TAO, Hive,  
2 Switchboard), Special Master Garrie finds that the burden of searching and producing from such  
3 systems outweighs any probative value the data may have. Issues regarding searching and producing  
4 from systems other than TAO, Hive, and Switchboard have been resolved, and the parties have removed  
5 all other systems from their proposals.

### 6 **ORDER**

7 29. No later than July 13, 2022, Facebook is to produce user objects and associations to those  
8 objects in the TAO system for each Named Plaintiff, along with the TAO schema for such data.  
9 Facebook is also to produce Named Plaintiffs' "set permissions" (audience controls on a post),  
10 including any audience controls on individual pieces of content, from TAO.

11 30. No later than August 1, 2022, Facebook is to produce Named Plaintiff data from the  
12 Switchboard system.

13 31. No later than August 8, 2022, Facebook is to produce the following types of Named Plaintiff  
14 data in Hive regardless of whether it appears in the DYI files: off-platform activity, ad interests, ad  
15 click data, ad impressions data, and custom audience data. Facebook will also provide the names of the  
16 tables from which the Hive data described above will be produced, how Facebook identified the tables,  
17 and the schema for such data

18 32. No later than July 25, 2022, Facebook is to produce the schema (column names) for all  
19 11,051 Hive tables identified in Exhibit B to Facebook's April 11, 2022 submission and the schema for  
20 all of the 137 preserved Hive tables that contain user identifiers.

21 33. No later than July 18, 2022, Facebook is to produce updated privacy settings for each  
22 Named Plaintiff as set out in Facebook's May 30, 2022 Letter.

23 34. No later than July 8, 2022, Facebook is to produce data regarding interactions that friends  
24 of the Named Plaintiffs had with businesses/apps using Facebook Login (without identifying which  
25 friend interacted with each business/app).


26 35. No later than August 7, 2022, each party is to provide to each other and Special Master Garrie  
27 a list of 250 Hive tables to be searched for Named Plaintiff data. Three business days after the 500  
28 tables have been identified, Facebook will submit a statement identifying which tables contain data in

### **AMENDED ORDER REGARDING PRODUCTION OF NAMED PLAINTIFF DATA**

1 cold storage and estimate the amount of data in cold storage for each table. Special Master Garrie will  
2 set a briefing schedule for proposals regarding searching cold storage at this time. In parallel with the  
3 cold storage analysis, Facebook is to search data in warm storage for the 500 Hive tables for data  
4 associated with the Named Plaintiffs and produce such data on a rolling weekly basis. Once Named  
5 Plaintiff data has been produced from all 500 Hive tables, Facebook is to determine what percentage of  
6 this data has not been produced from other sources and to submit a statement to this effect. Once  
7 Facebook has provided the results of this analysis Special Master Garrie will determine the appropriate  
8 next steps concerning Hive data.

9  
10  
11  
12  
13 IT IS SO ORDERED.

14  
15 Friday, July 1, 2022

16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
  
Daniel Garrie  
Discovery Special Master

# Exhibit A

Lesley E. Weaver (SBN 191305)  
BLEICHMAR FONTI & AULD LLP  
555 12th Street, Suite 1600  
Oakland, CA 94607  
Tel.: (415) 445-4003  
Fax: (415) 445-4020  
lweaver@bfalaw.com

Derek W. Loeser (admitted *pro hac vice*)  
KELLER ROHRBACK L.L.P.  
1201 Third Avenue, Suite 3200  
Seattle, WA 98101  
Tel.: (206) 623-1900  
Fax: (206) 623-3384  
dloeser@kellerrohrback.com

*Plaintiffs' Co-Lead Counsel*

*Additional counsel listed on signature page*

**UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA**

IN RE: FACEBOOK, INC. CONSUMER  
PRIVACY USER PROFILE LITIGATION

MDL No. 2843  
Case No. 18-md-02843-VC

This document relates to:

ALL ACTIONS

**PLAINTIFFS' SECOND SET OF  
REQUESTS FOR PRODUCTION TO  
DEFENDANT FACEBOOK, INC.**

Judge: Hon. Vince Chhabria  
Courtroom: 4, 17th Floor

PROPOUNDING PARTY: Plaintiffs

RESPONDING PARTY: Facebook

SET NUMBER: Two (2)

Plaintiffs hereby propound the following requests for production of documents to Defendant Facebook, Inc. (“Facebook”), pursuant to Federal Rules of Civil Procedure 26 and 34, and request that Facebook produce the documents and electronically-stored information set forth herein within thirty (30) days of service of these requests, at Bleichmar Fonti & Auld LLP, 555 12th Street, Suite 1600, Oakland, CA 94607.

### **INSTRUCTIONS**

1. You shall respond to these requests for the production of documents in a manner consistent with the Federal Rules of Civil Procedure and the following instructions:

2. In responding to each document request, furnish all responsive documents available at the time of production, including documents in your possession, custody or control, and in the possession, custody or control of your agents, employees, partners, representatives, subsidiaries, affiliates, investigators, or by your attorneys or their agents, employees or investigators.

3. If any otherwise responsive document was, but is no longer, in existence or in your possession, custody or control, identify the type of information contained in the document, its current or last known custodian, the location/address of such document, the identity of all persons having knowledge or who had knowledge of the document and describe in full the circumstances surrounding its disposition from your possession or control.

4. This is a continuing request for the production of documents and requires supplemental responses as provided for in the Federal Rules of Civil Procedure. If, after making your initial production, you (or any other persons acting on your behalf) obtain or become aware

of any further documents responsive to any document request, you are required to produce such additional documents to plaintiffs. Each supplemental response shall be served on plaintiffs no later than thirty days after the discovery of the further information.

5. You shall produce the original of each document described below or, if the original is not in your custody, then a copy thereof, and in any event, all non-identical copies which differ from the original or from the other copies produced for any reason, including, without limitation, the making of notes thereon.

6. Documents shall be produced as kept in the regular course of business together with the original folders, binders, boxes or other containers in which they were maintained.

7. All documents or things that respond in whole or in part to any portion of these requests are to be produced in their entirety, including attachments and their enclosures.

8. Documents attached to each other should not be separated.

9. Documents not otherwise responsive to any particular document request shall be produced if such documents mention, discuss, refer to, or explain the documents called for by any document request, or if such documents are attached to documents called for by any document request.

10. Documents shall be produced in such fashion as to identify the custodian of each document.

11. Identify the source of each document produced, by identifying: (a) all of the person(s) who possessed the document; (b) the positions or titles of any such individuals; and (c) all of the divisions and departments where each document was located. If you are unable to determine the individual(s) who possessed the document, identify the department and division where the document was located when produced.

12. If you claim any form of privilege, whether based on statute or otherwise, as a ground for not producing any document, state the following:

- a. The date of the document;
- b. The name, the present or last known home and business address, the telephone numbers, the title (or position), and the occupation of those individuals who prepared, produced, reproduced or who were recipients of said document;
- c. A description of the document sufficient to identify it without revealing the information for which the privilege is claimed;
- d. The nature of the privilege asserted;
- e. The factual basis upon which you claim any such privilege;
- f. The location of the document; and
- g. The custodian of the document.

13. To the extent you object to any document request, you must provide specific responses as to what portion of the request you object to and state expressly why you will not respond to such request in sufficient detail to permit the Court to determine the validity of the objection. Responsive documents to which your objection does not apply should be produced.

14. If you claim that all or any part of any document request, the Definitions, or Instructions is vague or ambiguous, please identify the specific language you consider vague or ambiguous and state the interpretation of the language in question you used to frame your response.

15. Each document requested herein is to be produced in its entirety and without deletion or excision, regardless of whether you consider the entire document to be relevant or responsive to any document request. If you have removed, excised or deleted any portion of a



document, stamp the word “REDACTED” on each page of the document that you have redacted. Redactions should be included on the privilege log described in Instruction No. 13, above.

16. One copy of each document should be produced. A document that varies in any way from the original or from any other copy, including drafts or a document with handwritten notations or deletions constitutes a separate document and must be produced, whether or not the original is in your possession, custody or control. Color (*i.e.*, not black and white) originals should be produced in color. If any identical copy cannot be produced for any reason (*e.g.*, faint writing, erasures, etc.), produce the original.

17. Indicate the origin of each document and number each document with consecutive Bates numbers.

### **DEFINITIONS**

Unless otherwise stated, the terms set forth below are defined as follows and shall be used in construing the meaning of these requests for the production of documents.

1. The use of the singular shall be deemed to include the plural, and the use of one gender shall include all others, as appropriate, in the context.

2. The present tense of a verb includes its past tense, and vice versa.

3. “And” and “or” are to be construed conjunctively and disjunctively, as necessary, to bring within the scope of this request for production all responses that might otherwise be construed to be outside its scope.

4. “Any” and “all” mean each and every.

5. “App” means an interactive software application developed to utilize the core technologies of the Facebook social networking platform.

6. “App Developer Investigation” or “ADI” means (as described in paragraph seven of the Chen Declaration) Facebook’s investigation to determine “whether there has been misuse

of data in violation of Facebook’s policies and associated legal liabilities, in connection with the first version of the [Facebook] Platform.”

7. “Apps Others Use” means the setting used to prevent the disclosure of personal information to third party App Developers through Facebook’s API, as described in paragraphs 366 to 368 of the FAC.

8. “App Settings” means settings that a User can alter or accept to limit Third Parties from accessing or obtaining Users’ Content and Information, including Apps Others Use, Granular Data Permissions, Platform Opt Out, and the like.

9. “Chen Declaration” means the Declaration of Stacy Chen in Support of Respondent’s Opposition to the Attorney General’s Petition, *Attorney General Maura Healy v. Facebook, Inc.*, No. 1984CV02597-BFS-1 (Mass. Super Ct., Suffolk Cty.).

10. “Communication” means the transmittal (in the form of facts, ideas, thoughts, opinions, data, inquiries or otherwise) and includes, but is not limited to, correspondence, memoranda, reports, presentations, face-to-face conversations, telephone conversations, text messages, instant messages, messages sent on Facebook Messenger, voice messages, negotiations, agreements, inquiries, understandings, meetings, letters, notes, telegrams, mail, electronic mail or email, and postings of any type.

11. “Computer System” or “Computer Systems” include(s), but is not limited to, any server (whether physical or virtual), desktop computer, tablet computer, point of sale system, smart phone, cellular telephone, networking equipment, internet site, intranet site, and the software programs, applications, scripts, operating systems, or databases used to control, access, store, add, delete, or modify any information stored on any of the foregoing non-exclusive list.

12. “Content and Information” refers to the definition in footnote 2 of the FAC, referring to “content” and “information” as Facebook’s Statements of Rights and

Responsibilities have defined those terms. In brief, Facebook has generally used “information” to mean facts and other information about Users, including the actions they take, and “content” to mean anything Users post on Facebook that would not be included in the definition of “information.” Content and Information also includes both personally identifiable content and information and anonymized content and information that is capable of being de-anonymized. *See* FAC ¶¶ 223-224. Content and Information includes data that identifies, relates to, describes, is capable of being associated with, or could reasonably be linked, directly or indirectly, with a particular User, including:

- a. Identifiers such as a real name, alias, postal address, unique personal identifier, online identifier, Internet Protocol address, email address, account name, social security number, driver’s license number, passport number, or other similar identifiers.
- b. Characteristics of protected classifications under California or federal law.
- c. Commercial information, including records of personal property, products or services purchased, obtained, or considered, or other purchasing or consuming histories or tendencies.
- d. Biometric information.
- e. Internet or other electronic network activity information, including, but not limited to, browsing history, search history, and information regarding a consumer’s interaction with an Internet Web site, application, or advertisement.
- f. Geolocation data.
- g. Audio, electronic, visual, thermal, olfactory, or similar information.
- h. Professional or employment-related information.

- i. Education information, defined as information that is not publicly available personally identifiable information as defined in the Family Educational Rights and Privacy Act (20 U.S.C. section 1232g, 34 C.F.R. Part 99).
- j. Inferences drawn from any of the information identified in this paragraph to create a profile, dossier, or similar collection of information about a consumer reflecting the consumer's preferences, characteristics, psychological trends, predispositions, behavior, attitudes, intelligence, abilities, and aptitudes.

13. "Document" or "Documents" is defined to include any Document, ESI, or Electronic Media stored in any medium, and is synonymous in meaning and equal in scope to the usage of this term in Federal Rule of Civil Procedure 34(a)(1)(A), including, but not limited to, electronic or computerized data compilations, Communications, electronic chats, instant messaging, documents created through Workplace by Facebook, encrypted or self-destructing messages, messages sent via Facebook messenger, email Communications, other electronically stored information from Personal computers, sound recordings, photographs, and hard copy Documents maintained in your Personal files.

14. "Electronic Media" means any magnetic, optical, or other storage media device used to record ESI including but not limited to computer memory, hard disks, floppy disks, flash memory devices, CDs, DVDs, Blu-ray discs, cloud storage (*e.g.*, DropBox, Box, OneDrive, or SharePoint), tablet computers (*e.g.*, iPad, Kindle, Nook, or Samsung Galaxy), cellular or smart phones (*e.g.*, BlackBerry, iPhone, or Samsung Galaxy), personal digital assistants, magnetic tapes of all types, or any other means for digital storage and/or transmittal.

15. "Electronically Stored Information" or "ESI" means information that is stored in Electronic Media, regardless of the media or whether it is in the original format in which it was created, and that is retrievable in perceivable form and includes, but is not limited to, metadata,

system data, deleted data, fragmented data, data pertaining to or maintained in Apps, database contents, and computer code.

16. “FAC” refers to the First Amended Consolidated Complaint filed February 22, 2019, ECF No. 257.

17. “Facebook,” “Defendant,” “You,” or “Your” shall mean Facebook, Inc. and any of its executives, directors, officers, employees, partners, members, representatives, agents (including attorneys, accountants, consultants, investment advisors or bankers), and any other Person purporting to act on its behalf. In the case of business entities, these defined terms include parents, subsidiaries, affiliates, predecessor entities, successor entities, these defined terms include parents, subsidiaries, affiliates, predecessor entities, successor entities, divisions, departments, groups, acquired entities and/or related entities or any other entity acting or purporting to act on its behalf.

18. “FTC Consent Order” shall refer to the July 27, 2012 Federal Trade Commission Consent Order in *In the Matter of Facebook, Inc.*, No. C-4365.

19. “Granular Data Permissions” means the setting through which the User accessing an App may limit the categories of Content and Information an App Developer may collect.

20. “Identify,” with respect to Documents, means to give, to the extent known, the (a) type of Document; (b) general subject matter; (c) date of the Document; (d) author(s); (e) addressee(s); and (f) recipient(s).

21. “Including” means “including but not limited to,” or “including, without limitation.” Any examples which follow these phrases are set forth to clarify the request, definition or instruction but not to limit the request.

22. “Internal Policy” or “Internal Policies” mean any formal or informal policy, procedure, rule, guideline, collaborative document, directive, instruction, or practice, whether written or unwritten, that You expect Your employees to follow in performing their jobs.

23. “Misuse of Data,” when used as a capitalized phrase, means the use by an App of a User’s Content or Information that was broader or different than the use of that content or information only in connection with the person that gave the permission to the App to access such User’s Content or Information.

24. “Named Plaintiffs” means Steven Akins, Jason Ariciu, Samuel Armstrong, Anthony Bell, Bridgett Burk, Brendan Carr, John Doe, Terry Fischer, Shelly Forman, Paige Grays, Mary Beth Grisi, Tabielle Holsinger, Taunna Lee Johnson, Olivia Johnston, Tyler King, Ashley Kmieciak, William Lloyd, Gretchen Maxwell, Scott McDonnell, Ian Miller, Jordan O’Hara, Bridget Peters, Kimberly Robertson, Scott Schinder, Cheryl Senko, Dustin Short, Tonya Smith, Mitchell Staggs, Charnae Tutt, Barbara Vance-Guerbe, and Juliana Watson.

25. “Person” or “Persons” means any natural Person or any business, legal or governmental entity or association.

26. “Platform” refers to the services, tools, and products provided by Facebook to third parties to create their own applications and services that access data in Facebook.

27. “Platform Opt Out” means the setting a User may access to choose that his or her Content and information is not accessed or obtained by any Apps or websites on Facebook’s Platform.

28. “Privacy Controls” means the audience selectors that control what information in a User’s profile can be viewed by other Users, and includes Profile Privacy Settings, Profile Privacy Controls, Publisher Privacy Controls, and the like.

29. “Relating to,” “relate to,” “referring to,” “refer to,” “reflecting,” “reflect,” “concerning,” or “concern” means all Documents which comprise, explicitly or implicitly refer to, were reviewed in conjunction with, or were created, generated or maintained as a result of the subject matter of the request, including, but not limited to, all Documents which reflect, record, memorialize, embody, discuss, evaluate, consider, review or report on the subject matter of the request.

30. “Third Parties” include the following:

- a. Apps, App Developers, Whitelisted Apps, and Business Partners, as those terms are used in the FAC;
- b. Any person that develops an application, software experience, game, or website that accesses Content and Information from Facebook’s API or other Facebook software; and
- c. Any person with which Facebook has or had an integration partnership.

31. “User(s)” means individuals who maintain a Facebook account and can generally access the typical Facebook experience through website or mobile applications.

32. Capitalized terms and acronyms not specifically defined herein have the same definition as in the FAC.

### **RELEVANT TIME PERIOD**

The relevant time period for each Document Request is January 1, 2007 through the present (the “Relevant Time Period”), unless otherwise specifically indicated. Each Document Request shall be interpreted to include all documents and information that relate to the Relevant Time Period or otherwise specified period, even if such documents or information were prepared or published outside of the Relevant Time Period or otherwise specified period. If a document prepared before or after this period is necessary for a correct or complete understanding of any

document covered by a request, you must produce the earlier or subsequent document as well. If any document is undated and the date of its preparation cannot be determined, the document shall be produced if otherwise responsive to the production request.

### **DOCUMENT REQUESTS**

#### **REQUEST FOR PRODUCTION NO. 6**

All Documents provided to or received from any governmental entity or regulator in the United States and United Kingdom in response to any formal or informal inquiry or investigation relating to whether Users' Content and Information was accessed or obtained by any Third Parties without proper consent or authorization, including but not limited to all inquiries or investigations arising out of the Cambridge Analytica Scandal, the FTC Consent Order, and any inquiry or investigation related to the settlement agreement with the FTC announced on July 24, 2019.

#### **REQUEST FOR PRODUCTION NO. 7**

All organizational charts, personnel directories, or other documents sufficient to show Your organizational structure, including:

- (a) the identity of subsidiaries, affiliates, and joint ventures, and your ownership interest, control of, or participation in any subsidiary or affiliate or joint venture related to agreements, engineering, access, use, transmission, receipt, collection or analysis of Facebook Users' Content and Information by Third Parties;
- (b) the organization of any division, department, unit or subdivision of your company that has responsibilities relating to agreements, engineering, access, use, transmission, receipt, collection or analysis of Users' Content and Information by Third Parties; and
- (c) the names, titles, job descriptions, and employment periods for your present and former employees who has or had responsibilities relating to agreements, engineering, access,



use, transmission, receipt, collection or analysis of Users' Content and Information by Third Parties; and

(d) the names, titles, job descriptions, and employment periods of Your present or former directors, officers, or senior managers, as well as any secretaries or administrative assistants assigned to these directors, officers, or senior managers.

**REQUEST FOR PRODUCTION NO. 8**

All versions (including each updated or amended version thereof) of Facebook's "Platform Policies," which have been called the "Developer Principles and Policies," the "Platform Guidelines," or the "Developer Terms of Service" (collectively, the "Platform Policies").

**REQUEST FOR PRODUCTION NO. 9**

All Documents relating to each of the Named Plaintiffs, including but not limited to all Content and Information collected about each of them or gained from business relationships or any other source.

**REQUEST FOR PRODUCTION NO. 10**

For each of the Named Plaintiffs, Documents sufficient to show the categories of Content and Information Facebook collects, tracks, and maintains about them.

**REQUEST FOR PRODUCTION NO. 11**

Documents sufficient to identify all Third Parties to which Facebook granted access to Named Plaintiffs' Content and Information, what categories of Content and Information Facebook granted access to, how Facebook allowed these Third Parties to access the Named Plaintiffs' Content and Information, and the business purpose of all such access.

**REQUEST FOR PRODUCTION NO. 12**

Documents relating to any partnerships or agreements Facebook entered into with Third Parties for access to Named Plaintiffs' Content and Information.

**REQUEST FOR PRODUCTION NO. 13**

For all Third Parties to which Facebook granted access to Named Plaintiffs' Content and Information, Documents sufficient to show any use by Third Parties of such Content and Information not in connection with the User that granted the permission to the Third Party or inconsistent with Facebook's agreement with that Third Party.

**REQUEST FOR PRODUCTION NO. 14**

Documents sufficient to show the monetary or retail value of each named Plaintiff's Content and Information to Facebook, updated to reflect whenever Facebook's terms of service changed, including the calculation of revenue earned by Facebook for each Named Plaintiff based upon bartering or selling access to such Named Plaintiff's Content and Information.

**REQUEST FOR PRODUCTION NO. 15**

Documents sufficient to show the money or any other thing of value, including but not limited to money or any other thing of value paid in exchange for targeted advertising, that Facebook received in exchange for each Named Plaintiff's Content and Information, which entities paid Facebook, and when such payments were made.

**REQUEST FOR PRODUCTION NO. 16**

Documents sufficient to show the monetary or retail value of Users' Content and Information to Facebook, including all monthly, quarterly, and annual financial reporting relating to same, and including but not limited to the calculation of average revenue per user, any changes to such monetary or retail value relating to changes to Facebook's terms of service, and any financial reporting of Content and Information as an asset.

**REQUEST FOR PRODUCTION NO. 17**

All Documents relating to Facebook’s assessment of the monetary or retail value of Users’ Content and Information to Users (as distinct from value to Facebook), including analyses for providing compensation to Users for their Content and Information, including but not limited to Users compensated in connection with the Onavo or Research app.

**REQUEST FOR PRODUCTION NO. 18**

All Documents that have been transmitted to Users by Facebook relating to whether Users’ Content and Information was accessed or obtained by Third Parties.

**REQUEST FOR PRODUCTION NO. 19**

All Documents supporting the escalation of those Apps escalated to Phase Two of ADI for Enhanced Examination and/or Phase Three of ADI for Enforcement and designated as follows in the Chen Declaration ¶ 34:

(d) each [A]pp to which a request for information was sent; (e) each [A]pp for which an interview was sought with the developer; (f) each [A]pp for which a remote or onsite audit was requested to be conducted; (g) each [A]pp for which actual misuse was found and identification of that misuse; (h) each [A]pp that was banned for actual misuse; and (i) each [A]pp that was banned for failing to cooperate with Facebook’s investigation.

Facebook has described identification of these Apps as non-privileged and has already produced it to the Massachusetts Attorney General’s Office. *See* Chen Declaration ¶ 35.

**REQUEST FOR PRODUCTION NO. 20**

The list of Apps that Facebook provided to the Massachusetts Attorney General’s Office and that the Chen Declaration ¶ 35 describes as “the subject of external actions or

communications with third parties, including the growing list of Apps Facebook has suspended as part of the [ADI], whether because of policy violations or because of their refusal to cooperate with Facebook's investigation."

**REQUEST FOR PRODUCTION NO. 21**

Communications between Facebook and Third Parties relating to the ADI, including but not limited to Communications that Facebook provided to the Massachusetts Attorney General's Office. *See* Chen Declaration ¶ 37.

**REQUEST FOR PRODUCTION NO. 22**

All "Privacy Risk Assessment[s]," and notes or agenda relating to Facebook's "focused subject-matter-specific meetings," "focused subject-matter-specific discussions," "weekly intra- and inter-team meetings," and "Privacy Summit[s]," as detailed in "Facebook's Privacy Program Overview" included in any PricewaterhouseCoopers LLP ("PwC") assessment report prepared pursuant to the FTC Consent Order.

**REQUEST FOR PRODUCTION NO. 23**

Unredacted versions and Documents in support of the assessment reports, including the Initial Assessment Report and Biennial Reports, prepared by PwC pursuant to the FTC Consent Order.

**REQUEST FOR PRODUCTION NO. 24**

Documents sufficient to identify all Third Parties to which Facebook granted access to Users' Content and Information not generally available through Platform pursuant to partnerships or agreements between Facebook and those Third Parties.

**REQUEST FOR PRODUCTION NO. 25**

All Documents relating to agreements or partnerships described in Request No. 24.

#### **REQUEST FOR PRODUCTION NO. 26**

For each of the Third Parties that Facebook entered into partnerships or agreements with as described in Request No. 24, Documents sufficient to identify:

- The fields, kinds, or categories of Content and Information that were accessed or obtained by such Third Parties;
- How each such Third Party accessed or obtained the Content and Information of Users;
- How each such Third Party used the Content and Information accessed or obtained;
- Where the Content and Information obtained by such Third Parties currently resides and who has access to it.

#### **REQUEST FOR PRODUCTION NO. 27**

Documents sufficient to show all forms and formats in which Facebook transmitted to Third Parties information concerning Users' liking, viewing, retrieving, or otherwise requesting or obtaining videos on, using, or by means of the Facebook Platform.

#### **REQUEST FOR PRODUCTION NO. 28**

All Documents relating to Internal Policies by Facebook on the monitoring of Third Parties' compliance with Facebook's Platform Policy, Data Policy, or SRR.

#### **REQUEST FOR PRODUCTION NO. 29**

All Documents relating to Internal Policies by Facebook on the enforcement of Facebook's Platform Policy, Data Policy, or SRR against Third Parties.

#### **REQUEST FOR PRODUCTION NO. 30**

All Documents relating to measures and controls, including proposed measures and controls, put in place by Facebook to prevent Third Parties from violating Facebook's Platform Policy, Data Policy, or SRR.

**REQUEST FOR PRODUCTION NO. 31**

All Documents relating to Facebook's audits, inquiries, and investigations of Third Parties investigating compliance with any provisions of Facebook's Platform Policy, Data Policy, or SRR regarding the access, use, transmission, receipt, collection and analysis of Users' Content and Information on and off the Platform.

**REQUEST FOR PRODUCTION NO. 32**

All Documents Concerning Misuse of Data, including investigations, examinations, inquiries, or audits—or Communications regarding such investigations, examinations, inquiries, or audits—regarding Misuse of Data prior to the deprecation of Graph API v.1.0.

**REQUEST FOR PRODUCTION NO. 33**

Documents sufficient to show the notice that Facebook provided to Users regarding modifications to Facebook's SRR or Data Policy, and all Communications related thereto.

**REQUEST FOR PRODUCTION NO. 34**

All Documents relating to the conditioning of Third Parties' access to Users' Content and Information on the purchase of Mobile App Install Ads, payment of Content and Information in-kind (referred internally as Reciprocity or Data Reciprocity), or other payment.

**REQUEST FOR PRODUCTION NO. 35**

Documents relating to the manner in which a Facebook User could control how his or her data was shared through their Privacy Controls and App Settings throughout the Relevant Time Period, including but not limited to screenshots of the Facebook website and the Facebook mobile application.

**REQUEST FOR PRODUCTION NO. 36**

All Documents concerning User testing, evaluation and analysis of Facebook's Privacy Controls and App Settings during the Relevant Time Period, including but not limited to design documents, correspondence, analyses, and reports.

Dated: November 25, 2019

Respectfully submitted,

KELLER ROHRBACK L.L.P.

BLEICHMAR FONTI & AULD LLP

By: /s/ Derek W. Loeser  
Derek W. Loeser

By: /s/ Lesley E. Weaver  
Lesley E. Weaver

Derek W. Loeser (admitted *pro hac vice*)  
Lynn Lincoln Sarko (admitted *pro hac vice*)  
Gretchen Freeman Cappio (admitted *pro hac vice*)  
Cari Campen Laufenberg (admitted *pro hac vice*)  
Benjamin Gould (SBN 250630)  
1201 Third Avenue, Suite 3200  
Seattle, WA 98101  
Tel.: (206) 623-1900  
Fax: (206) 623-3384  
dloeser@kellerrohrback.com  
lsarko@kellerrohrback.com  
gcappio@kellerrohrback.com  
claufenberg@kellerrohrback.com  
bgould@kellerrohrback.com

Lesley E. Weaver (SBN 191305)  
Anne K. Davis (SBN 267909)  
Joshua D. Samra (SBN 313050)  
555 12th Street, Suite 1600  
Oakland, CA 94607  
Tel.: (415) 445-4003  
Fax: (415) 445-4020  
lweaver@bfalaw.com  
adavis@bfalaw.com  
jsamra@bfalaw.com

Christopher Springer (SBN 291180)  
801 Garden Street, Suite 301  
Santa Barbara, CA 93101  
Tel.: (805) 456-1496  
Fax: (805) 456-1497  
cspringer@kellerrohrback.com

*Plaintiffs' Co-Lead Counsel*



# Exhibit B

## Exhibit A

What info is available?	What is it?	Where can I find it?
About Me	Information you added to the <b>About</b> section of your timeline like relationships, work, education, where you live and more. It includes any updates or changes you made in the past and what is currently in the <b>About</b> section of your timeline.	Downloaded Info
Account Status History	The dates when your account was reactivated, deactivated, disabled or deleted.	Downloaded Info
Active Sessions	All stored active sessions, including date, time, device, IP address, machine cookie and browser information.	Downloaded Info
Address	Your current address or any past addresses you had on your account.	Downloaded Info
Ads	Ads you've recently viewed.	Downloaded Info
Ads Clicked	Dates, times and titles of ads clicked (limited retention period).	Downloaded Info
Ad Topics	A list of topics that you may be targeted against based on your stated likes, interests and other data you put in your timeline.	Downloaded Info
Advertising ID	The unique advertising identification numbers provided by your mobile device. These numbers are used to show you ads on the apps you use on your device.	Downloaded Info
Alternate Name	Any alternate names you have on your account (example: a maiden name or a nickname).	Downloaded Info
Apps	All of the apps you have added.	Downloaded Info
Articles	Articles you've recently read.	Downloaded Info
Autofill Information	Information you've provided, such as your address, that is used to pre-fill messages when you contact a business through Messenger.	Downloaded Info
Chat	A history of the conversations you've had on Facebook Chat (a complete history is available directly from your messages inbox).	Downloaded Info

Source: *What categories of my Facebook data are available to me?*, <https://www.facebook.com/help/930396167085762>, Table 2, *Information you can download using the Download Your Information tool* (last visited Sept. 18, 2020).

Chat Rules	Chat Rules you've accepted.	Downloaded Info
Check-ins	The places you've checked into.	Downloaded Info
Currency	Your preferred currency on Facebook. If you use Facebook Payments, this will be used to display prices and charge your credit cards.	Downloaded Info
Current City	The city you added to the <b>About</b> section of your timeline.	Downloaded Info
Date of Birth	The date you added to Birthday in the <b>About</b> section of your timeline.	Downloaded Info
Dating	The number of times you've recently visited the Dating section of Facebook.	Downloaded Info
Device ID	The unique identification numbers provided by the devices you use to log into Facebook.	Downloaded Info
Device Locale	The country and language from which you're accessing Facebook as determined by the devices you're using.	Downloaded Info
Education	Any information you added to Education field in the About section of your timeline.	Downloaded Info
Emails	Email addresses added to your account (even those you may have removed).	Downloaded Info
Email Address Verifications	A history of when you've verified your email address.	Downloaded Info
Events	Events you've joined or been invited to.	Downloaded Info
Event Contacts You've Blocked	People you've blocked from inviting you to events.	Downloaded Info
Event Interactions	The number of times you've recently visited the Events section of Facebook.	Downloaded Info
Events Visited	Event pages you've recently visited.	Downloaded Info
Facebook Live Videos	Live videos you've recently watched.	Downloaded Info
Facebook Watch Topics for Recommendations	A collection of topics that is used to show you relevant videos in the Facebook Watch tab. The topics are	Downloaded Info

	based on your previous interaction history with things like links, videos, photos and Pages you've liked.	
Facial Recognition Data	A unique number based on a comparison of the photos you're tagged in. We use this data to help others tag you in photos.	Downloaded Info
Family	Friends you've indicated are family members.	Downloaded Info
Favorite Quotes	Information you've added to the Favorite Quotes section of the <b>About</b> section of your timeline.	Downloaded Info
Followers	A list of people who follow you.	Downloaded Info
Friends	A list of your friends.	Downloaded Info
Friend Requests	Pending, sent and received friend requests.	Downloaded Info
Friends You See Less	Friends whose activity you've chosen to see less of on Facebook.	Downloaded Info
Fundraisers	Fundraisers you've recently viewed.	Downloaded Info
Gender	The gender you added to the <b>About</b> section of your timeline.	Downloaded Info
Groups	A list of groups you belong to on Facebook.	Downloaded Info
Group Interactions	The number of times you've interacted with Groups on Facebook.	Downloaded Info
Groups Visited	Groups you've recently visited.	Downloaded Info
Hometown	The place you added to hometown in the <b>About</b> section of your timeline.	Downloaded Info
ID	A copy of the ID you submitted to confirm your identity and to help improve our automated systems for detecting fake IDs and related abuse.	<a href="#">Personal Data Request</a>
Instant Games	Instant Games you've played.	Downloaded Info
IP Address Activity	Your recent activity from specific IP addresses.	Downloaded Info

IP Address Message Activity	Your recent message activity from specific IP addresses.	Downloaded Info
IP Address Payment Activity	Your recent payment activity from specific IP addresses.	Downloaded Info
Language Settings	Your preferred language settings.	Downloaded Info
Last Location	Your most recent location determined by your device.	Downloaded Info
Linked Accounts	Accounts you've linked to your Portal.	Downloaded Info
Live Video Subscriptions	Scheduled Live videos you've subscribed to.	Downloaded Info
Logins	IP address, date and time associated with logins to your Facebook account.	Downloaded Info
Logouts	IP address, date and time associated with logouts from your Facebook account.	Downloaded Info
Marketplace Categories	Categories you've recently viewed.	Downloaded Info
Marketplace Interactions	Your recent interactions on Marketplace.	Downloaded Info
Marketplace Items	Items you've recently viewed.	Downloaded Info
Marketplace Services	Services you've recently viewed.	Downloaded Info
Matched Contacts	Contact information that may be associated with your account.	<a href="#">Personal Data Request</a>
Menu Items	Areas of Facebook you've recently accessed through the main menu.	Downloaded Info
Messages	Messages you've sent and received on Facebook. Note, if you've deleted a message it won't be included in your download as it has been deleted from your account.	Downloaded Info
Messenger Contacts You've Blocked	Contacts you've blocked on Messenger.	Downloaded Info
Milestone Notifications	Notifications about your activity milestones, such as the number of reactions on a post, you've received and dismissed.	Downloaded Info

Mobile Service Provider and Country Code	The service provider and country code associated with your phone number.	Downloaded Info
Name	The name on your Facebook account.	Downloaded Info
Name Changes	Any changes you've made to the original name you used when you signed up for Facebook.	Downloaded Info
News Feed Topics for Recommendations	A collection of topics that is used to show you relevant public posts in parts of your News Feed. The topics are based on your previous interaction history with things like links, videos, photos and Pages you've liked.	Downloaded Info
News Topics for Recommendations	A collection of topics that is used to show you relevant articles in the News tab. The topics are based on your previous interaction history with things like posts, videos, photos and Pages you've liked.	Downloaded Info
Notification ID	The identification numbers that we use to send you Facebook notifications on your device.	Downloaded Info
Page Notifications	Chat notifications you've dismissed from Pages you visit.	Downloaded Info
Page Visits	Pages you've recently visited.	Downloaded Info
Page Transparency Notices	A list of pages that you've received and dismissed notices from.	Downloaded Info
Pages You Admin	A list of pages you admin.	Downloaded Info
Pages You've Recommended	Pages you've recommended to others.	Downloaded Info
Pending Friend Requests	Pending, sent and received friend requests.	Downloaded Info
People	People and friends you've interacted with recently, including comments and reactions.	Downloaded Info
People Viewed	People you've recently viewed when new friends were suggested to you.	Downloaded Info
Phone Numbers	Mobile phone numbers you've added to your account, including verified mobile numbers you've added for security purposes.	Downloaded Info

Photos	Photos you've uploaded to your account.	Downloaded Info
Photo Effects	A list of the photo effects you've used.	Downloaded Info
Photos Metadata	Any metadata that is transmitted with your uploaded photos.	Downloaded Info
Platforms	Platforms you've used to log into Facebook, such as the Facebook app or a browser.	Downloaded Info
Pokes	A list of who's poked you and who you've poked. Poke content from our mobile poke app is not included because it's only available for a brief period of time. After the recipient has viewed the content it's permanently deleted from our systems.	Downloaded Info
Political Views	Any information you added to Political Views in the About section of timeline.	Downloaded Info
Preferred Language for Videos	The preferred language for videos as determined by videos you've previously viewed.	Downloaded Info
Previously Removed Contacts	Friends you've recently removed but added back.	Downloaded Info
Primary Location	Your primary location is determined by information we use to support Facebook Products, such as the current city you entered on your profile and your device connection information.	Downloaded Info
Profile Visits	People whose profiles you've recently visited.	Downloaded Info
Recent Activities	Actions you've taken and interactions you've recently had.	Downloaded Info
Recently Visited	Videos and shows you've recently visited.	Downloaded Info
Record Details	Details included in some administrative records.	Downloaded Info
Registration Date	The date you joined Facebook.	Downloaded Info
Religious Views	The current information you added to Religious Views in the <b>About</b> section of your timeline.	Downloaded Info
Removed Friends	People you've removed as friends.	Downloaded Info

Saved Post Reminders	Reminders you've received after you've saved a post.	Downloaded Info
Screen Names	The screen names you've added to your account, and the service they're associated with. You can also see if they're hidden or visible on your account.	Downloaded Info
Secret Conversations	A list of the times you've used Secret Conversations in Messenger.	Downloaded Info
Secret Conversations You've Reported	A list of the secret conversations you've reported to Facebook.	Downloaded Info
See First	Profiles and Pages you've recently chosen to see first in your News Feed.	Downloaded Info
See Less	Profiles and Pages you've recently chosen to see less of in your News Feed.	Downloaded Info
Selected Language	The language you've selected to use Facebook in.	Downloaded Info
Session Type	Your current active session types.	Downloaded Info
Show Pages	A list of the Show Pages you've viewed and the videos you've watched from them.	Downloaded Info
Shows	A list of the individual videos you've watched.	Downloaded Info
Spoken Languages	The languages you added to Spoken Languages in the <b>About</b> section of your timeline.	Downloaded Info
Status Updates	Any status updates you've posted.	Downloaded Info
Time Spent	The amount of time you've spent watching videos from a Show Page.	Downloaded Info
Time Viewed	The amount of an individual video you've watched.	Downloaded Info
Timezone	The timezone you've selected.	Downloaded Info
Work	Any current information you've added to Work in the <b>About</b> section of your timeline.	Downloaded Info
Videos	Videos you've posted to your timeline.	Downloaded Info



Video Creator Pages	Video creator Pages you've recently viewed.	Downloaded Info
Videos You've Removed	Videos you've removed from your Watch list.	Downloaded Info
Your Facebook Activity	A history of when you've accessed Facebook.	Downloaded Info
Your Pinned Posts	Posts you've pinned on your timeline.	Downloaded Info

# Exhibit C

Pages 1 - 22

UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

Before The Honorable Jacqueline Scott Corley, Magistrate Judge

IN RE FACEBOOK, INC. CONSUMER )  
PRIVACY USER PROFILE )  
LITIGATION. ) NO. 18-MD-02843 VC (JSC)  
\_\_\_\_\_ )

San Francisco, California  
Friday, August 14, 2020

**TRANSCRIPT OF REMOTE VIDEOCONFERENCE PROCEEDINGS**

**APPEARANCES VIA ZOOM:**

For Plaintiffs:

KELLER ROHRBACK LLP  
1201 Third Avenue, Suite 3200  
Seattle, Washington 98101  
BY: DEREK W. LOESER, ATTORNEY AT LAW  
DAVID J. KO, ATTORNEY AT LAW

BLEICHMAR, FONTI & AULD LLP  
555 - 12th Street, Suite 1600  
Oakland, California 94607  
BY: LESLEY E. WEAVER, ATTORNEY AT LAW  
ANNE K. DAVIS, ATTORNEY AT LAW  
ANGELICA M. ORNELAS, ATTORNEY AT LAW  
MATTHEW P. MONTGOMERY, ATTORNEY AT LAW

For Defendants:

GIBSON, DUNN & CRUTCHER LLP  
200 Park Avenue  
New York, New York 10166-0193  
BY: ORIN SNYDER, ATTORNEY AT LAW

**(APPEARANCES VIA ZOOM CONTINUED ON FOLLOWING PAGE)**

Reported Remotely By: Ana M. Dub, RDR, CRR, CCRR, CRG, CCG  
Official Reporter, CSR No. 7445

**APPEARANCES VIA ZOOM:** (CONTINUED)

For Defendants:

GIBSON, DUNN & CRUTCHER LLP  
Trammell Crow Center  
2001 Ross Avenue, Suite 2100  
Dallas, Texas 75201

**BY: RUSSELL H. FALCONER, ATTORNEY AT LAW**

GIBSON, DUNN & CRUTCHER LLP  
333 S. Grand Avenue  
Los Angeles, California 90071

**BY: DEBORAH L. STEIN, ATTORNEY AT LAW**

GIBSON, DUNN & CRUTCHER LLP  
1881 Page Mill Road  
Palo Alto, California 94304

**BY: MARTIE KUTSCHER CLARK, ATTORNEY AT LAW**

Friday - August 14, 2020

8:27 a.m.

P R O C E E D I N G S

---o0o---

**THE CLERK:** We're a minute early, but court is now in session. Let's see. Calling Civil Action 18-MD-2843, In Re Facebook, Inc. Consumer Privacy User Profile Litigation.

Counsel, starting with plaintiff, can you please state your appearance.

**MS. WEAVER:** Sure. This is Lesley Weaver of Blakemar Fonti & Auld. With me is Anne Davis and Angelica Ornelas.

And I see that Matt Montgomery actually is not -- he should be with us. So he should probably be elevated. I apologize. I missed him before. Don't tell him.

**MR. LOESER:** Good morning. You have Derek Loeser from Keller Rohrbach.

**THE COURT:** Good morning.

**MR. KO:** Good morning, Your Honor. Nice to see you again. David Ko, Keller Rohrbach, also on behalf of plaintiffs.

**THE COURT:** Good morning.

And here comes Mr. Montgomery. He's here.

All right. And for Facebook?

**MR. SNYDER:** Good morning, Judge. It's Orin Snyder from Gibson Dunn with my colleagues, Deb Stein, Martie Kutscher Clark, and Russ Falconer.

1           **THE COURT:** Good morning.

2           Okay. Thank you for your statement.

3           Let's see. It sounds like there are not too many things  
4 to discuss. Let's just start.

5           The search terms you're working on, I will just make this  
6 observation. I do think it would be unreasonable to insist  
7 that all terms apply to all custodians. That just can't be  
8 right. People have different positions. So I give you that  
9 guidance in working on that.

10          Now, with respect to the data about plaintiffs, let's go  
11 through. And why don't plaintiffs tell us what is the data  
12 that you're missing that you think is relevant. So one thing  
13 you've identified is the data about what data about the  
14 plaintiffs was shared with advertisers. Is that correct?

15           **MS. WEAVER:** That is correct in general terms,  
16 Your Honor. Basically, what has been produced to us is  
17 user-facing data through an Access Your Account tool, for the  
18 most part.

19          Now, I want you to know that we have reviewed all of the  
20 plaintiffs' data with more than one pass-through. We've done  
21 targeted searches. We've had 18 people, and more at times,  
22 going through the documents. So we're pretty familiar with  
23 what's there.

24          There are two problems that we have. The first is that  
25 Your Honor ordered us last -- two weeks ago to discuss

1 precisely what has been produced and precisely what is the data  
2 that is being withheld.

3 And we -- in the course of our meet-and-confer sessions,  
4 Facebook did not identify the examples that they put in their  
5 statement. We didn't discuss those. So once again, we are  
6 getting information the first time in the statement.

7 And it would have been better if we had discussed it,  
8 because when we look at those documents -- we've looked at them  
9 before -- they are not what we're seeking. And the reason that  
10 they're not -- and if you look, there's an example of one of  
11 them they gave us. The content is missing. So there's an  
12 event that says one of the users went to a website, but the  
13 content of what they did on the site is stripped away.

14 And our experts say, you know, what did you put in your  
15 shopping cart? What did you access? How long were you on it?

16 And that data is also married to GPS data --

17 **THE COURT:** Okay. I have the statement --

18 **MS. WEAVER:** Yeah.

19 **THE COURT:** -- in front of me.

20 **MS. WEAVER:** Yes.

21 **THE COURT:** Can you put me to the page and the Bates  
22 number?

23 **MS. WEAVER:** The Bates number of the document -- hang  
24 on.

25 **THE COURT:** Well, first, the page of the statement so

1 I know where to go.

2 **MS. WEAVER:** That is going to be harder for me.  
3 I think it's page 6. The Bates number -- and I'm going to  
4 ask -- Anne, if you can help me, it's 01037245.

5 **THE COURT:** Don't see that. It's redacted  
6 information?

7 **MS. WEAVER:** Some of the information was redacted,  
8 yes. But this information we can discuss in the hearing, if  
9 that is --

10 **THE COURT:** No, no. I understand. We can -- I'm not  
11 worried --

12 **MS. WEAVER:** Yeah.

13 **THE COURT:** -- about that.  
14 I'm just trying to find it. I don't see it.

15 **MS. WEAVER:** Yeah. Hang on just a moment.

16 **THE COURT:** Maybe the sentence at the first page of  
17 the --

18 **MS. WEAVER:** Yeah, I'm actually looking -- I  
19 apologize. I'm looking for the actual statement. I have too  
20 many things open on my laptop.

21 But for all of the documents that they've identified,  
22 Your Honor, these are PDFs that reflect some activity.

23 **THE COURT:** I just want to start with -- I want to  
24 start with --

25 **MS. WEAVER:** Fine. Okay. So if you go to page 5 of



1 the statement and if we look at, for example, where it says  
2 "Ms. Tutt reviewed content on Amtrak.com," it doesn't tell us  
3 what the content is or it doesn't tell us --

4 **THE COURT:** Okay. Or the --

5 **MS. WEAVER:** -- what they did.

6 **THE COURT:** -- other one, that Ms. Tutt viewed content  
7 on a news site and --

8 **MS. WEAVER:** Right. And it doesn't --

9 **THE COURT:** -- tell you what the content is.

10 **MS. WEAVER:** -- tell us what they do.

11 **THE COURT:** Let me ask Facebook.

12 Do you have that content?

13 **MR. SNYDER:** Mr. Falconer, I think, will address this.

14 **MR. FALCONER:** Good morning, Your Honor. Russ  
15 Falconer for Facebook.

16 Our understanding is there is some machine-readable data  
17 in some cases that might reflect the off-Facebook activity that  
18 Ms. Weaver is describing in a kind of raw, disaggregated way.  
19 That information is not associated with the plaintiff's account  
20 in the way that the user-created, user-shared content and  
21 information is associated with a user account.

22 And so I hear -- I don't know -- confusion and frustration  
23 from Ms. Weaver that they feel like they don't understand what  
24 we've produced.

25 The Court ordered us to, you know, be as clear as we can

1 on named plaintiffs' data, what has been produced and what has  
2 been withheld. And what we've tried to do is say that we've  
3 produced all content information that the plaintiffs share on  
4 Facebook and then some of the other categories of information  
5 that we identified in our statement; so device information,  
6 geolocation information, certain other information that is  
7 associated with their account. And we have been -- I think  
8 we've tried to be clear; and if we failed in this, we  
9 apologize.

10 There is other -- there's Facebook-generated information,  
11 information generated by third parties, information received  
12 from third parties. We have not represented that that is  
13 comprehensively included in our production.

14 What we have produced are Facebook analytics, third-party  
15 data, off-Facebook activity, anything like that that is  
16 associated with a user's account.

17 And so that's -- I think the point of departure between  
18 the parties right now is maybe the level of generality with  
19 which we have described what we have not produced. But  
20 that's -- we've tried to be as clear about the, sort of, large  
21 buckets that are not included in the named plaintiff data we've  
22 produced to date.

23 **THE COURT:** So, for example, when you say Ms. Weaver  
24 said, as you said, that the plaintiff viewed content on  
25 Amtrak.com, are you saying you don't have any way of

1 identifying what that content is that she viewed at that  
2 particular time, even though you were able to say she viewed  
3 that website at that time?

4 **MR. FALCONER:** I think for an individual plaintiff on  
5 an individual website, if it was just that question -- could we  
6 tell for one of the named plaintiffs what specific content she  
7 viewed on the Amtrak website? -- if it was, you know, ten years  
8 ago or seven years ago, probably not. If it was a year ago,  
9 maybe. That data may or may not have been associated with --

10 **THE COURT:** Well, if it was this year --

11 **MR. FALCONER:** Yeah.

12 **THE COURT:** -- with that particular --

13 **MR. FALCONER:** Sure.

14 **THE COURT:** -- data this year.

15 **MR. FALCONER:** The answer is it's possible. There may  
16 be some website-specific data about that named plaintiff; there  
17 may not be. There's some --

18 **THE COURT:** Okay. And so you haven't searched for it,  
19 or you're withholding it, or -- I guess, why hasn't it been  
20 produced?

21 **MR. FALCONER:** So as we understood the Court's  
22 mandate or, sort of, the Court's --

23 **THE COURT:** No, no, no. I'm just asking.

24 **MR. FALCONER:** Oh.

25 **THE COURT:** I'm just asking.

1           **MR. FALCONER:** Because the reason for that is that  
2 just to find it for one named plaintiff would be like a  
3 multiweek endeavor, if not longer. And the reason for that is  
4 that -- let's take the Amtrak example.

5           With this off-Facebook activity data, the tables and the  
6 database where the data is stored, you know, they've been  
7 explained to us like each one of them is a book. And the book  
8 is organized by topic. The topic that the book is organized by  
9 is the advertiser. It's Amtrak; it's not the named plaintiff.

10          So for every Facebook advertiser there's a book. Right?  
11 There's a table that has some data for advertisement, website  
12 activity, that kind of thing.

13          So to gather the information for one named plaintiff on  
14 Amtrak, that, we could probably do. To gather the data for one  
15 named plaintiff on every advertiser on every off-Facebook  
16 activity that has ever happened, just for one named plaintiff,  
17 we have to go into each of those books individually and look  
18 for that one named plaintiff, and then we'd have to do it for  
19 each of the other 23 named plaintiffs.

20          So that's the reason why we have not undergone that to  
21 date.

22           **THE COURT:** I understand that. So have you identified  
23 every instance that you have that the plaintiff viewed content  
24 on some website, whatever it is?

25           **MR. FALCONER:** Every instance where Facebook has been

1 able to associate that off-Facebook activity with a named  
2 plaintiff's account. Sometimes they can't make the connection.  
3 But where it's connected, we've identified it. That's included  
4 in the production.

5 **THE COURT:** I assume that for this privacy case --  
6 right? -- some content is obviously more private than other  
7 content and the plaintiffs may not necessarily need or want.  
8 They need exemplars. Right? And there is a standing argument  
9 that you guys are maintaining that they have to defeat and  
10 damages and all that. There are particular instances. Right?  
11 So there may be particular instances where you then have to go  
12 do that.

13 In other words, if it's the data that was shared, which is  
14 sort of at the heart of the case, you're probably going to have  
15 to do some work on that. Whether it's every instance, probably  
16 not; but certainly certain instances.

17 Now, plaintiffs, it sounds like, have a template of where  
18 to start. It may not be Amtrak, but it may be the next one  
19 there. Right?

20 **MR. FALCONER:** Your Honor, could I be heard on that?

21 **MS. WEAVER:** Well, may I --

22 **MR. FALCONER:** Or, go ahead.

23 **MS. WEAVER:** I would like to respond.

24 So what we're talking about right now and what they've  
25 produced is, there's a tool so users can download data. And

1 even in what they're downloading, there is content missing.

2 But there's another whole bucket of data that they haven't  
3 identified to us that is responsive, and that's the first step.  
4 We need the identification of the fields of the data that they  
5 collect through their third-party relationships, whether it's  
6 apps or websites, et cetera. And it is this database that  
7 Facebook searches using algorithms to target the users.

8 What they've given us is sort of the window dressing of  
9 the platform activity, and I've identified for you that  
10 something is missing even from that.

11 But there is -- and, Your Honor, we've talked to our  
12 experts; and maybe it's better to have experts talk or put in a  
13 declaration because I can tell you, their position will be that  
14 this is, quote/unquote, not associated with the users but that  
15 doesn't make sense.

16 There is an event ID, because the reason Facebook is  
17 collecting it in the first place is to target people with the  
18 data. So there is a way to go back and find -- and I agree  
19 with Mr. Falconer that this data set will be immense. And that  
20 is the scope of the case. And that's why we said only for the  
21 24 because --

22 **THE COURT:** I'm just going to --

23 **MS. WEAVER:** Yeah.

24 **THE COURT:** -- tell you guys, I think maybe you need  
25 to think about a special master.

1       There's just no --

2           **MS. WEAVER:** Yes.

3           **THE COURT:** I don't have the time or the patience or  
4 the expertise to wade through any of this, like the nuance that  
5 you're getting into. So I don't know what to do.

6           **MS. STEIN:** Your Honor, may I be heard for a moment?

7       So I think the good news on, sort of, your reaction to  
8 this is that this exercise was really about, sort of,  
9 identifying categories so that we could have a conversation  
10 about what's required in this case, because there is a whole  
11 lot of information being sought here that has absolutely  
12 nothing to do with the issues that are being litigated in this  
13 case.

14          **THE COURT:** No. I understand that argument. I don't  
15 even know how to figure out what it is that we're even talking  
16 about.

17          **MS. STEIN:** Right.

18          **MS. WEAVER:** So Facebook --

19          **MS. STEIN:** So, Your Honor, what's being --

20          **MS. WEAVER:** Could I --

21          **MS. STEIN:** -- talked about right now is what's called  
22 off-Facebook activity. And that off-Facebook activity has no  
23 relationship to the issues that the dismissal order said are  
24 viable right now and that are not stayed. The order of  
25 dismissal --

1           **THE COURT:** No. I read that. I read it. I  
2 understand.

3           **MS. STEIN:** Okay. Good.

4           **THE COURT:** So this --

5           **MS. STEIN:** And so the off-Facebook activity --

6           **THE COURT:** -- this has been previewed -- just, can I  
7 finish?

8           **MS. STEIN:** I'm sorry, Your Honor.

9           **THE COURT:** Because I'm really losing patience with  
10 this case.

11           This has been previewed for a while. So what I was hoping  
12 to do is you guys could just tee up what that data is so I can  
13 rule if it's discoverable or not.

14           I don't even know how to get to that point.

15           **MR. SNYDER:** Your Honor, I think there's a very  
16 easy --

17           **MS. WEAVER:** If I could, I was waiting.

18           Your Honor, we would like them to identify what they're  
19 withholding. That's it.

20           **THE COURT:** But that's a chicken-and-egg problem.  
21 That's a chicken-and-egg problem. And I'm not sure -- and see,  
22 this is the problem I'm having. You said you've now reviewed  
23 it all. What is missing? You've identified --

24           **MS. WEAVER:** So I'll give you examples. There are no  
25 examples --



1           **THE COURT:** You did.

2           **MS. WEAVER:** Okay.

3           **THE COURT:** No. I'm going to let Mr. Snyder talk.

4           **MS. WEAVER:** Fine.

5           **MR. SNYDER:** Your Honor, I share your frustration, and  
6 I think this is very easy.

7           For example, on advertisement, we have gone, I think as  
8 indicated in our statement, above and beyond the call of duty  
9 because we didn't really want to just say, "We're not giving  
10 you what advertisements you reviewed or ads that you've clicked  
11 on, even though it's outside the scope of the case."

12          This case --

13          **THE COURT:** No, no. That's an argument. Please,  
14 let's try not to argue.

15          **MR. SNYDER:** Right.

16          **THE COURT:** I'm going to decide that at some point.

17          **MR. SNYDER:** Okay. So what I would --

18          **THE COURT:** Just --

19                   (Simultaneous cross-talk.)

20          **THE COURT:** -- that.

21          **MR. SNYDER:** What I would respectfully suggest is, we  
22 can, Your Honor, tee it up for you in a very simple way,  
23 because Judge Chhabria's order is very clear about what's in  
24 and what's out. And then each side can succinctly,  
25 efficiently, and clearly make their arguments about what is in

1 and what's out. And it's not going to be difficult,  
2 Your Honor. I think it's pretty clear.

3 I agree, on this call, people using terminology --  
4 "on-platform," "off-platform" -- it all sounds like  
5 gobbledegook. I think there's a very clear, efficient, and  
6 efficacious way for us to tee this up in a short statement to  
7 Your Honor; and Your Honor can rule on it, if Your Honor wants  
8 more argument on it, without us having these dueling  
9 Zoom/Hollywood Squares, you know, arguments about what's in and  
10 what's out that's not going to really lead to any fair ruling.

11 **THE COURT:** This is what I need to ask Ms. Weaver, is:  
12 Do you know what it is that you want or that you believe exists  
13 that you don't have?

14 **MS. WEAVER:** Yes.

15 **THE COURT:** You do. Okay.

16 **MS. WEAVER:** More or less. We don't know what form  
17 they keep it in or how they keep it. It is this data set that  
18 they mine, yes.

19 **THE COURT:** Okay. So is there any reason why, then,  
20 we can't adjudicate that dispute as discoverability?

21 **MS. WEAVER:** We can --

22 **MR. SNYDER:** I think we can --

23 **MS. WEAVER:** -- adjudicate that, Your Honor.

24 **THE COURT:** We can? Okay.

25 **MR. SNYDER:** We can and we should.

1           **THE COURT:** All right.

2           **MR. SNYDER:** And I think we can do it very simply  
3 without a lot of drama or complication.

4           **THE COURT:** So that's what --

5           **MR. FALCONER:** Your Honor --

6           **THE COURT:** -- I want you to do, then, on this,  
7 I think.

8           And, I mean, it doesn't have to be the joint letter brief,  
9 whatever. I mean, it's a big issue. It kind of goes to the  
10 heart of the case. So I want you to have the ability. You're  
11 going to probably need your experts to some extent -- at least  
12 plaintiffs -- to be involved with it.

13           And I probably want four briefs. Right? Whoever goes  
14 first, second, first, second, so that there's -- my guess is  
15 it's not till we get to the second two briefs that we'll really  
16 be able to meet there. That just seems to be the process that  
17 we need to do.

18           So you guys work it out, how that's going to be presented.  
19 I'm not giving you any limits at all. You only have the limit  
20 of my time and attention span. So just keep that in mind.

21                               (Laughter.)

22           **MS. WEAVER:** And how much time, Your Honor, would you  
23 like between briefs and the hearing? What kind of timing --

24           **THE COURT:** We'll put a hearing. I'll figure it out.

25           **MS. WEAVER:** Okay.

1           **THE COURT:** I mean, to be honest, I'm just swamped at  
2 the moment.

3           **MS. WEAVER:** I know.

4           **THE COURT:** So, but you get it to us. We'll get  
5 through it. And we will set it for hearing. I think it's  
6 important to have an oral --

7           **MR. LOESER:** And, Your Honor, if I could be heard for  
8 one quick minute on one --

9           **THE COURT:** Yes.

10          **MR. LOESER:** This is Derek Loeser.  
11 -- just, process point.

12          Where we stand right now, we generally think we know  
13 what's missing, and we can describe it in our briefs.

14          Facebook obviously has specific knowledge about what's  
15 missing. And so because they haven't identified specifically  
16 what they're withholding, I really think it would be improper  
17 for them to argue in their brief that we haven't been specific  
18 enough with what we're seeking. If that is going to be their  
19 argument in their brief, then they should comply with your last  
20 order, which was to identify specifically what they're  
21 withholding.

22          But that's the only --

23          **THE COURT:** Yeah. No, I understood. So that's why  
24 I'm doing four briefs.

25          And in the meantime, you should be talking and really

1 trying to narrow. It is in both sides' interest to have it  
2 teed up as accurately as possible for me to decide. Otherwise,  
3 I'm going to make a wrong decision one way or the other because  
4 I won't understand.

5 **MR. SNYDER:** And, Your Honor, it's in everyone's  
6 interest to have you not be frustrated with us, which I  
7 understand and I think your frustration is well-placed, one.

8 Two, we want Your Honor to continue to preside over  
9 discovery; and we would, I think, lose a lot if we had to start  
10 fresh with a special master.

11 And mindful of that, we're going to work to narrow the  
12 issues. Maybe we can even eliminate them. And we have a lot  
13 of other work to do in the meantime. So however long  
14 Your Honor needs, we're going to obviously abide and respect  
15 that, and we're not going to, you know, ask you to turn around  
16 a ruling.

17 There's a lot we have to do on search terms and privilege  
18 logs and ADI protocols. So there's a ton of work for us to do  
19 while Your Honor takes -- you know, takes the time necessary to  
20 adjudicate this issue, which is ripe now.

21 **THE COURT:** Yeah. Just don't put a hearing date.  
22 I'll pick it. So that's not a problem.

23 **MR. LOESER:** The only thing I would add to that,  
24 Your Honor, is that we would like you to be very frustrated  
25 with Orin all the time, but not with us.

1 (Laughter.)

2 **THE COURT:** Well, this week has not been -- I've been  
3 frustrated a lot, and I apologize for that.

4 **MR. SNYDER:** Don't apologize.

5 **MS. WEAVER:** It's tough times.

6 **THE COURT:** There's a lot. There's just a lot,  
7 scheduling.

8 **MR. SNYDER:** Yes, Your Honor.

9 **THE COURT:** Okay. So, which leads me to my next  
10 point, which is the joint statement -- okay? -- which is, you  
11 all are extremely talented, experienced lawyers. If you can't  
12 figure out a way, a process for this statement to work -- it's  
13 really, actually, for you. Right? The statement is a great  
14 way of assessing where we are, what our disputes are,  
15 crystallizing it. It's for you more than me, quite honestly.  
16 And if you guys can't figure out together a way to do that,  
17 then we've got to go back to zero and start over. I mean, this  
18 should be the easy part.

19 So I'm not going to tell you how to do that joint  
20 statement. The only thing I'm going to tell you is I want it  
21 however -- what is -- just even one day, I give you, right,  
22 before this? I take it upon myself; I will make time to read  
23 it the night before or early the morning before. That's my  
24 only deadline. You guys work it out. Whatever works best for  
25 you and gets it. But the point is, it should really try to

1 crystallize it.

2 My own view is -- and with other cases -- is that -- at  
3 least with discovery disputes, is if you do time for a reply as  
4 opposed to changing what you've already said, that tends to  
5 work better. But I'm not ruling at all. I want you guys to  
6 come up with it. It's, frankly, below my pay grade to have to  
7 tell you how to do it.

8 (Laughter.)

9 **MR. LOESER:** We hear that loud and clear, Your Honor,  
10 and we will keep talking to Facebook about it.

11 We just think that it would be really useful for everyone  
12 here, including for you, if people talk about things that they  
13 put in their statements before it's submitted to the Court.  
14 And so that's our mission in trying to come up with a better  
15 way to do this. That's what we're trying to accomplish.

16 **THE COURT:** Maybe you could do a statement, a draft,  
17 and then you talk about what's in the draft. Right? So then  
18 you know what's in there before you -- I don't know, but that  
19 would --

20 **MR. LOESER:** Yeah. We'll figure it out.

21 **THE COURT:** Yes. I know you guys can figure it out  
22 because you're all outstanding lawyers. That's why you're on  
23 this case.

24 Okay. So then we need to pick our next date. How about  
25 we push it out three weeks, to September 3rd?

1           **MR. LOESER:** I think that's the 749th day of March;  
2 so, sounds great.

3                               (Laughter.)

4           **MS. WEAVER:** That's fine, Your Honor.

5           **MR. SNYDER:** And two months before Election Day,  
6 assuming the post offices --

7           **MS. WEAVER:** There is one.

8           **MR. SNYDER:** -- assuming the post offices and the  
9 polling places aren't shut down permanently.

10          **THE COURT:** All right. Okay.

11          **MR. LOESER:** Don't depress us, Orin.

12          **THE COURT:** I apologize for having to lecture a little  
13 bit, but to be honest, you guys can do better. I know you can.  
14 I know you can. I have tremendous respect for all of you.

15          Okay. Great. I look forward to our next conference.  
16 It'll be September 3rd at 8:30 a.m.

17          **MR. SNYDER:** Thank you, Judge.

18          **MS. WEAVER:** Thank you, Your Honor.

19          **MR. SNYDER:** Thank you for everything you're doing.  
20 Appreciate it.

21          **THE CLERK:** Court's adjourned.

22                               (Proceedings adjourned at 8:51 a.m.)

23                               ---o0o---

24

25



**CERTIFICATE OF REPORTER**

I certify that the foregoing is a correct transcript  
from the record of proceedings in the above-entitled matter.

DATE: Saturday, August 15, 2020

*Ana M. Dub*

Ana M. Dub, CSR No. 7445, RDR, CRR, CCRR, CRG, CCG  
Official Reporter, U.S. District Court

# Exhibit D

Derek W. Loeser (admitted *pro hac vice*)  
KELLER ROHRBACK L.L.P.  
1201 Third Avenue, Suite 3200  
Seattle, WA 98101  
Tel.: (206) 623-1900  
Fax: (206) 623-3384  
dloeser@kellerrohrback.com

Lesley E. Weaver (SBN 191305)  
BLEICHMAR FONTI & AULD LLP  
555 12th Street, Suite 1600  
Oakland, CA 94607  
Tel.: (415) 445-4003  
Fax: (415) 445-4020  
lweaver@bfalaw.com

*Plaintiffs' Co-Lead Counsel*

*Additional counsel listed on signature page*

**UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA**

IN RE: FACEBOOK, INC. CONSUMER  
PRIVACY USER PROFILE LITIGATION

MDL No. 2843  
Case No. 18-md-02843-VC

This document relates to:

ALL ACTIONS

**PLAINTIFFS' OPPOSITION TO  
DEFENDANT FACEBOOK, INC.'S  
REQUEST TO ENFORCE THE PARTIAL  
STAY OF DISCOVERY IN PRETRIAL  
ORDER NO. 20 AND CROSS-MOTION  
TO COMPEL DISCOVERY RELATED  
TO REQUESTS FOR PRODUCTION  
NOS. 9 THROUGH 13**

Judges: Hon. Vince Chhabria  
Hon. Jacqueline S. Corley  
Courtroom: 4, 17th Floor

## TABLE OF CONTENTS

I.	INTRODUCTION .....	1
II.	ARGUMENT .....	2
A.	The Order does not limit discovery to users’ platform activity.....	2
B.	The discovery requests at issue and Facebook’s response .....	6
C.	Relevant sensitive information is not limited to platform activity, but also includes sensitive information Facebook derives and collects from business partners, app developers, apps, and other sources. ....	7
1.	User data includes, in Facebook’s words, “native, appended and behavioral data” that Facebook collects from business partners, apps and other activity.....	8
2.	Internal documents confirm that Facebook’s description of data “associated” with users is misleading.....	11
D.	Facebook has not established that the burden of producing the data relating to ten Plaintiffs is disproportional to the needs of this case. ....	12
III.	CONCLUSION .....	14

## **TABLE OF AUTHORITIES**

	<b>Page(s)</b>
<b>Cases</b>	
<i>Doe v. Gangland Prods., Inc.</i> , 730 F.3d 946 (9th Cir. 2013) .....	10
<i>Harris v. Best Buy Stores, L.P.</i> , No. 3:15-cv-00657-HSG (KAW), 2016 WL 6024556 (N.D. Cal. Oct. 14, 2016).....	13
<i>Polaris Innovations Ltd. v. Kingston Tech. Co.</i> , No. CV1600300CJCRAOX, 2017 WL 3275615 (C.D. Cal. Feb. 14, 2017).....	14
<i>Shulman v. Grp. W. Prods., Inc.</i> , 18 Cal.4th 200 (1998).....	10
<i>SPS Techs., LLC v. Briles Aerospace, Inc.</i> , No. CV 18-9536 MWF, 2020 WL 4341717 (C.D. Cal. June 25, 2020) .....	14
<i>Sullivan v. Personalized Media Commc'ns, LLC</i> , No. 16-MC-80183-MEJ, 2016 WL 5109994 (N.D. Cal. Sept. 21, 2016) .....	14
<b>Statutes</b>	
18 U.S.C. § 2702(a) .....	3, 4, 5, 10
18 U.S.C. § 2710(b)(2).....	3, 4, 5, 10
<b>Other Authorities</b>	
Fed. R. Civ. P. 26(b)(1).....	13
Local Rule 5-1(i)(3) .....	16

## I. INTRODUCTION

Facebook does not want Plaintiffs to obtain discovery showing the full breadth of its wrongful disclosure of its users' sensitive information. Accordingly, Facebook seeks to limit discovery in this case to a single category of improperly shared information: users' activity on the Facebook platform. The sensitive information that Facebook collects and shares with third parties is much more extensive than this. It collects users' sensitive information from a variety of sources—including from third parties—then pools the information with user-posted activity and generates additional information from the full data set it accumulates. It then shares this information about users and their friends with third parties. All of this information, including who has access to it and how it is used, is relevant to Plaintiffs' claims.

As a result, there are at least three compelling reasons that Facebook's motion should be denied and Plaintiffs' cross-motion to compel production of documents responsive to Requests for Production ("RFPs") Nos. 9 through 13<sup>1</sup> should be granted.

*First*, contrary to Facebook's tortured reading of Pretrial Order No. 20 ("Order" or "PTO 20"), Dkt. No. 298, the Court did not limit discovery in this case only to information regarding user activity on Facebook. While that information—and Facebook's subsequent disclosure of it—is of course relevant, that is not the only type of sensitive information relevant to Plaintiffs' claims or the four categories of wrongdoing recognized by the Order.

*Second*, the universe of data Facebook collects and shares about users is also not limited to user activity on Facebook, but instead consists of a sea of information obtained from a wide variety of sources, including from business partners, app developers, apps, and other third parties. Indeed, as Facebook's own documents show, it collects information about users far beyond what Facebook has produced in this case. And discovery produced to date further confirms that Facebook not only collects this information, but links it to users and shares it with third parties—putting to rest Facebook's nonsensical suggestions that Plaintiffs have failed to articulate what additional evidence exists or that Facebook cannot "associate" certain data with a

---

<sup>1</sup> For details on these RFPs, see *infra* § II.B.

user.

*Third*, there is no justification for Facebook’s claims of undue burden. Such an argument should be accorded minimal weight in a case of this size and complexity involving a company whose business model is premised upon the collection and production of electronic information about billions of users. Facebook has come nowhere near meeting its burden of demonstrating why data regarding solely Named Plaintiffs—relative to the hundreds of millions of potential class members whose information is ultimately at issue in this case—is not proportional to the needs of the case. In fact, pursuant to the Court’s recent guidance regarding streamlining Plaintiffs’ discovery, Plaintiffs have reduced the number of individuals who will be class representatives to ten, down from the twenty-three. Plaintiffs only seek the discovery at issue here related to these ten Plaintiffs (for purposes of this motion, the “Named Plaintiffs.”)

## II. ARGUMENT

### A. The Order does not limit discovery to users’ platform activity.

PTO 20 does not directly address the question raised by Facebook in its motion—whether this case is limited to user activity on the Facebook platform or includes all the sensitive information about users that Facebook improperly shared with third parties. But the Order nowhere expressly limits the case to user activity. *Cf.* Mot.<sup>2</sup> at 1. Nor does it make sense to read the Order that way. That sort of limitation would conflict not only with claims and theories that the Order upheld, but also with the grounds on which they were allowed to proceed to discovery.

Facebook, under the guise of enforcing a discovery stay that was never issued in the first place, spends many pages straining to read the Order to limit discovery to data relating only to users’ on-platform activity. This provides a misleading picture of what the Order says and inaccurately ascribes to the Court a set of internally inconsistent views.

**1. The Order.** The Order summarizes its understanding of Plaintiffs’ claims in a two-sentence précis near the beginning: “Broadly speaking, this case is about whether Facebook

---

<sup>2</sup> Def. Facebook, Inc.’s Opening Brief in Supp. of Its Req. to Enforce the Partial Stay of Discovery in Pretrial Order No. 20 (“Mot.”), Dkt. No. 515.

acted unlawfully in making user information widely available to third parties. It's also about whether Facebook acted unlawfully in failing to do anything meaningful to prevent third parties from misusing the information they obtained." Order at 3. This description focuses on *Facebook's* unlawful disclosure of information about users and their friends to third parties—not on whether that information was originally posted, shared, or generated by users on the Facebook platform.

The Order then discusses the four categories of Facebook's wrongdoing. These categories are: (1) "[g]iving app developers access to sensitive user information"; (2) "[c]ontinued disclosure to whitelisted apps"; (3) "[s]haring sensitive user information with business partners"; and (4) "[f]ailure to restrict the use of sensitive information." Order at 6-9. These categories line up neatly with the earlier description of the action as alleging that "Facebook acted unlawfully in making user information widely available to third parties" (the first three categories) and that Facebook "fail[ed] to do anything meaningful to prevent third parties from misusing the information they obtained" (the fourth category). *Id.* at 3.

Using these four categories of wrongdoing as a framework, the Order analyzed whether Plaintiffs had standing to bring their claims and whether they stated valid claims. It ruled that Plaintiffs had standing because they alleged that their "sensitive information was disseminated to third parties in violation of their privacy." *Id.* at 14. It upheld nearly all of Plaintiffs' claims (e.g., three privacy-based tort claims under California law, a claim under the Stored Communications Act ("SCA"), a claim for breach of contract, and a claim for unjust enrichment) except to the extent they were based on the first category of wrongdoing, the disclosure of user information to app developers. *Id.* at 30-34, 38-41. It upheld in its entirety Plaintiffs' claim under the Video Privacy Protection Act ("VPPA"). *Id.* at 34-35. And it upheld Plaintiffs' claim for negligence, which was based on the fourth category of wrongdoing. *Id.* at 35-36.

**2. *The Order's rationale.*** Why did the Order conclude that Plaintiffs had standing and had stated valid claims? On these points, the Order is clear. Plaintiffs had standing because "their "sensitive information was disseminated to third parties in violation of their privacy." *Id.* at 14.



This reasoning focuses not on *where* the user information was originally generated—whether on the Facebook platform or off it—but on its nature (“sensitive”) and on what Facebook did with it (“disseminated” it “to third parties”).

Similarly, when discussing the claims, the Order focused not on the original provenance of the information about users, but on its nature and on what Facebook did with it. So, for example, the Order ruled that:

- Plaintiffs had stated valid privacy torts because Facebook had disseminated information that was “sensitive” and as to which Plaintiffs had a reasonable expectation of privacy. *Id.* at 30-33.
- Plaintiffs had stated a claim under the Stored Communications Act because Facebook had disseminated the content of their electronic communications and had not gained their consent to do so. *Id.* at 33-34.
- Plaintiffs had stated a claim under the Video Privacy Protection Act because Facebook had disseminated “information which identifies a person as having requested or obtained specific video materials or services,” *id.* at 34 (citation omitted), and Facebook qualified as a “video tape services provider” under the statute, *id.* at 35.

**3. “Sensitive information” is not defined by where Facebook collects that information.**

The Order repeatedly notes that Facebook shares “sensitive” user information without consent. Facebook pins its argument to this one word, maintaining that the Order “defined” sensitive user information to mean only information about what users post on Facebook, Mot. at 1, or users’ platform activity, Mot. at 8. But the common-sense meaning of “sensitive information” encompasses more than just what users did on the platform. Consider, for example, a Facebook user’s Amazon.com order for an over-the-counter contraceptive or another user’s entry of “alcoholic support group in Tower District, Fresno” into a search engine. “Sensitive information” also includes information that Facebook can *infer* from on-platform information—a category of information it has not produced. (Think of the inferences that Facebook can draw from weekly photographs of a user taken at M.D. Anderson Cancer Center.) Facebook’s objection that such information is categorically not “sensitive” is false.

It is true that when the Order gave examples of sensitive user information, the examples it

used concerned information generated on the Facebook platform. *E.g.*, Order at 1, 17. Nowhere, however, did the Order *define* or *limit* sensitive information to users’ platform activity only. And the Order’s reasoning certainly is not limited to such information. Rather, as noted above, Plaintiffs’ standing to bring their claims, and the validity of many of those claims, depends on the nature rather than the provenance of the information, and on whether Facebook shared that information with third parties. And, as Plaintiffs have learned through discovery, the sensitive information about users that Facebook collects and shares with business partners and app developers includes both information originally generated outside the Facebook platform and information derived from on- and off-platform activity.

It also is farfetched for Facebook to argue that the Order rules that *all* of Plaintiffs’ claims—including their federal statutory claims—rise or fall depending on whether the information that Facebook shared is “sensitive” in the sense of being embarrassing or deeply intimate. The validity of Plaintiffs’ claim under the SCA, for example, does not turn on how embarrassing or intimate the information is that Facebook shared, but on whether the shared information includes the contents of an electronic communication. 18 U.S.C. § 2702(a)(1). Similarly, Plaintiffs’ VPPA claim turns on whether the information that Facebook shared includes “information which identifies a person as having requested or obtained specific video materials or services from a video tape service provider.” *Id.* § 2710(a)(3). If, for example, Facebook collected and shared a user’s video-watching queue from a different platform, that would constitute a VPPA violation.

In sum, while the Order does not explicitly address the issue posed by Facebook’s motion, it certainly does not limit discovery in this case to on-platform user activity and reading it that way is inconsistent with the Court’s reasoning. It is also inconsistent with statements by the Court during the motion to dismiss hearing about the breadth of user data that is relevant to Plaintiffs’ claims:

For example, if – I’m a Facebook user. And, you know, I’m trying to assess the likelihood that my sensitive information got into the hands of third parties and, if so, how many third parties and, if so, what kinds of third parties. If I have a full

understanding of the third parties that had access to the information, and a full understanding of what type of information they had access to, and a full understanding of who they were, and what they – and what restrictions were placed on them, we then have a better understanding of what was likely to have happened to me.

Nov. 4, 2019 Tr. at 15:20-16:4. It is the “full understanding” referred to by the Court that Plaintiffs seek, and that Facebook refuses to allow.

Finally, this reading prevents Named Plaintiffs from discovering even the general policies and practices of Facebook governing the sharing of their sensitive information, policies and practices that are critical for this case. *See* Pretrial Order No. 30 at 2, Dkt. No. 347 (“[T]he best way to assess the merits and to determine whether class certification is appropriate is almost certainly to conduct discovery on Facebook’s general practices.”). Plaintiffs submit that Facebook’s exclusion of this information from discovery is not what the Order intended.

**4. The Order stayed claims, not discovery.** Plaintiffs organized their claims into three categories: prioritized claims, prioritized consumer protection act claims alleged in the alternative, and non-prioritized claims. First Am. Consolidated Compl. (“FACC”) at 317-411, Dkt. No. 257. The Order made the simple observation that “[a]ll other prioritized claims not addressed by this ruling will be stayed (effectively, relegated to non-prioritized status) and adjudicated, if necessary, at a later state in the proceedings with the other non-prioritized claims.” Order at 6. Facebook’s claim that this holding somehow imposed a stay of *discovery* is baffling. The Order does not, and does not purport to, stay discovery in any fashion.<sup>3</sup>

#### **B. The discovery requests at issue and Facebook’s response**

The present dispute arises from five discovery requests, each of which asks for data that Facebook possesses about Named Plaintiffs, the third parties that Facebook disclosed this data to, and the types of information that was disclosed to them. *See* Ex. A, Def. Facebook, Inc.’s Resps. & Objs. to Pls.’ Second Set of Reqs. for Produc. In particular, RFP No. 9 requests “[a]ll

---

<sup>3</sup> Even if it were, the Order observed that “[o]f course, dismissal of a subset of claims with prejudice does not preclude a plaintiff from seeking revival if discovery reveals a factual basis that justifies reconsideration[.]” Order at 37 n.21 (citations omitted).

Documents relating to each of the Named Plaintiffs, including but not limited to all Content and Information collected about each of them or gained from business relationships or any other source.”<sup>4</sup> *Id.* RFP No. 10 asks Facebook to produce, “[f]or each of the Named Plaintiffs, Documents sufficient to show the categories of Content and Information Facebook collects, tracks, and maintains about them.” *Id.* RFP Nos. 11-13 then request documents requesting Facebook to identify the third parties that were able to access this information, including the categories of data that were disclosed to them and how they accessed it. *Id.* Plaintiffs propounded these requests nearly one year ago in November 2019.

In response to these requests, Facebook produced information collected by the DYI (“Download Your Information”) tool. This limited tool allows downloads of some, but not all, information relating to users’ activity on the platform. And Facebook freely acknowledges that Plaintiffs can access this information themselves. *Id.* (“[A]ll Facebook users are free to download their DYI file if they wish.”). In addition to the DYI production, Facebook has produced an undefined category of “additional information associated with [users’] accounts” for each Plaintiff. Mot. at 6. But Facebook does not describe what the “additional information” is, likely because it is extremely limited—it consists solely of information users can access through their account in the form of their privacy settings and information reflecting user activity on Facebook. Critically, the form of production also obscures whether some of the activity was public or private. Thus, virtually all of Facebook’s 850,000-page production relating to the original Named Plaintiffs in this case was already accessible to Plaintiffs and tells only part of the story.

**C. Relevant sensitive information is not limited to platform activity, but also includes sensitive information Facebook derives and collects from business partners, app developers, apps, and other sources.**

Facebook acknowledges that it collects and shares substantial amounts of additional sensitive information about users beyond their platform activity. *See, e.g.*, Aug. 14, 2020 Hr’g

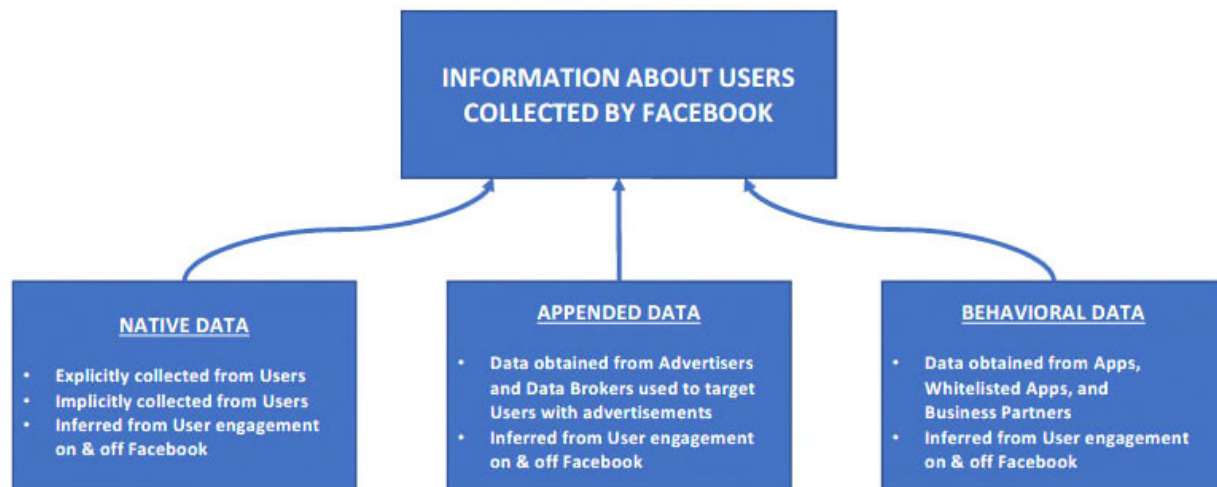
---

<sup>4</sup> The requests use the definition of “Content and Information” from Facebook’s Statement of Rights of Responsibilities—a definition that is not limited to on-platform data.

Tr. 8:10-13 (“[T]here’s Facebook-generated information, information generated by third parties, information received from third parties. We have not represented that that is comprehensively included in our production.”); *see also* Mot. at 10-15 (describing off-platform activity and internal analytics it has not produced). However, Facebook contends that this other information is not relevant to this case. This is false.

**1. User data includes, in Facebook’s words, “native, appended and behavioral data” that Facebook collects from business partners, apps and other activity.**

Facebook identifies three general categories of information it possesses about Facebook users: Native Data, Appended Data, and Behavioral Data. *See* Ex. B, FB-CA-MDL-00213424-439. Native Data includes public and private information Facebook collects from a user’s activity on Facebook such as a user’s profile info, posts, likes, shares, and location and device information. *Id.* Native Data also includes a small subset of information it can infer from user activity on and off Facebook, such as user interests and behaviors. *Id.* Appended and Behavioral Data, on the other hand, consists of information inferred from user activity on and off Facebook. Specifically, Appended Data is user information obtained from third parties, including advertisers and data brokers, about a user’s activities—*e.g.*, auto registration, retail and credit card purchase histories, as well as other “enhanced” customer databases. *Id.* Behavioral Data consists of substantial information related to users including their activity within third-party applications—*e.g.*, website browsing behavior, app installations, and purchases they make off the Facebook platform. *Id.* These three categories of information, and how Facebook obtains this information, are illustrated as follows:



While Native Data—the only type of data Facebook has partially produced about users so far—is important for this case, so too is Appended and Behavioral Data obtained from user “engagement” on and off the platform. Moreover, Facebook connects and integrates Native, Appended and Behavioral Data generated from on-platform activity with data obtained from third parties. *See* Ex. B, FB-CA-MDL-00213424-439 (identifying Facebook’s matching of data it receives directly from users with Appended Data and Behavioral Data received from third parties through “[REDACTED]”); *id.* at FB-CA-MDL-00213424 (describing Facebook’s [REDACTED] process as instances where [REDACTED]).

Critically—and contrary to Facebook’s suggestion that this data is irrelevant and duplicative of information it has already produced (Mot. at 14)—discovery confirms that Facebook shares this data with third parties. For example, a 2012 internal email between senior team members explains that “Facebook has information about users which can be helpful to applications, and *which we provide to applications when we deem appropriate[.]*” Ex. C, Email from Sam Lessin to Douglas Purdy (Aug. 30, 2012), at 3 (emphasis added), NBCNews (Nov. 6, 2019), <https://dataviz.nbcnews.com/projects/20191104-facebook-leaked-documents/assets/facebook-sealed-exhibits.pdf> at 716. The email also identifies the kind of data

that Facebook shares, including (1) “[a]ggregate data about the tastes, properties, etc. of a user’s friends,” (2) “[d]erived data about a user/facebook’s data/opinion of a user (probably location, account trust score, account age, etc.),” and (3) “[d]ata provided by third parties – information which third parties have contributed to the graph on behalf of a user.” *Id.*

Another internal document, [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED] These

documents make clear that Facebook collects sensitive user information in a variety of different ways and discloses it to third parties.

Facebook’s insistence that it need only produce on-platform Native Data makes even less sense when considering Plaintiffs’ claims. Plaintiffs’ statutory and common law claims are not limited to information generated from users’ activities on Facebook. For example, under the VPPA, Plaintiffs must prove that Facebook disclosed “personally identifiable information concerning any consumer” to “any person” absent written or informed consent. 18 U.S.C. § 2710(b)(2). Under the SCA, Plaintiffs must prove that Facebook “knowingly divulge[d] to any person or entity the contents of any communication” users did not intend for Facebook to divulge. 18 U.S.C. § 2702(a). The source of the information—that is, whether it was the result of on- or off- platform activity, gleaned directly from users’ posts, or inferred from them—is irrelevant. Disclosure of any of this information without consent is actionable.

Similarly, Plaintiffs’ Public Disclosure of Private Acts claim requires Plaintiffs to prove that Facebook disclosed a private fact about the plaintiff that is objectionable and offensive to a reasonable person. *Doe v. Gangland Prods., Inc.*, 730 F.3d 946, 958 (9th Cir. 2013). Likewise, Plaintiffs’ Intrusion into Private Affairs claim requires Plaintiffs to prove an intrusion by Facebook into a private matter that is highly offensive to a reasonable person. *Shulman v. Grp. W. Prods., Inc.*, 18 Cal.4th 200, 231 (1998). In order to prove these claims, Plaintiffs must



ascertain the private facts about them that Facebook is collecting and disclosing, whether they originate from platform activity or not.

Across many claims, the Order sustained Plaintiffs’ allegations about Facebook’s undisclosed data reciprocity programs with business partners. Plaintiffs are thus entitled to know what sensitive user data, of any type or source, Facebook shared with its business partners. Plaintiffs are further entitled to any data that Facebook received from its business partners in return, since the value of that data constitutes the benefit Facebook received in the transaction, a benefit that Plaintiffs are entitled to recover under, *inter alia*, the unjust enrichment claim that the Court sustained. Order at 41;<sup>5</sup> *see also* Order at 8 (noting the allegation that “Facebook and its [business] partners agreed to exchange information about users’ activities with each other”).

Facebook notes repeatedly that targeted advertising and psychographic marketing are not part of this case. *See, e.g.*, Mot. at 9. This argument misses the point. The question is not whether Facebook should or should not have engaged in targeted advertising and psychographic marketing. The question is whether, when doing so, Facebook shared sensitive user and friend information without consent. Plaintiffs are entitled to obtain the discovery necessary to substantiate the allegation that improper sharing has occurred in the context of these activities.

**2. Internal documents confirm that Facebook’s description of data “associated” with users is misleading.**

Facebook claims it has produced all data it possesses that is “associated” with Named Plaintiffs. That is, while it generated and collected reams of data about Named Plaintiffs, Facebook claims that most of that data, including Appended and Behavioral Data, is anonymized and cannot be connected to Named Plaintiffs. This is false.

Facebook explains that Appended and Behavioral Data cannot be associated with Plaintiffs’ Facebook accounts because such data is “disassociated from the user’s ID within 90

---

<sup>5</sup> Facebook’s position blocking discovery of what it possesses and shares is in tension with Facebook’s own discovery requests to Named Plaintiffs. Facebook’s Interrogatory No. 8 asks Plaintiffs to “Identify all entities other than Cambridge Analytica that You believe have “misused sensitive information from Your Facebook Account.” But Facebook itself will not identify with whom it shared that sensitive information, let alone what information it possesses.



days” (Mot. at 15). But, as confirmed by internal documents, what actually happens is that Facebook replaces the User ID, or “UID,” with a Random ID, or “Replacement ID” or “RID”—and *stores the connection between a UID and an RID* until users delete their accounts.<sup>6</sup> In other words, Facebook replaces the User ID with another identifier, yet the data is still directly linked to and associated with a given user. Indeed, the very purpose of collecting all of this data in the first place is to use it to target users and their friends. Facebook can readily connect any active user with Appended Data and Behavioral Data containing that user’s RID<sup>7</sup> and, thus, it can be easily searched for, identified, and copied, without “multiple operations . . . to try to determine whether any data in the analysis originated from a Named Plaintiff.” *Cf.* Mot. at 15.

Similarly, Appended Data without a RID can be connected to active users through “[h]ashed data matching.” *See* Ex. B. “Hashed data matching” is the process of matching different data sets through the hash values of unique identifiers. For instance, when an advertiser uploads a spreadsheet of Custom Audience data including hashed email addresses, Facebook can match this data to its users through the hashed email address field.

Thus, it simply is untrue that it would be “nearly impossible” to produce the “disassociated” data in this case for Named Plaintiffs. Mot. at 15. Facebook clearly has the ability to connect Named Plaintiffs’ user information through RIDs and hashed data matching, and should be ordered to do so in response to RFP Nos. 9-13.

**D. Facebook has not established that the burden of producing the data relating to ten Plaintiffs is disproportional to the needs of this case.**

Facebook also suggests that “the burdens of locating the additional information Plaintiffs seek would far exceed the needs of the case.” Mot. at 12. But the burden associated with producing the requested information is not undue; it is proportional to the needs of this complex

<sup>6</sup> Ex. E, PwC\_CPUP\_FB00030737-738.

<sup>7</sup> *Id.* at PwC\_CPUP\_FB00030738

case. In assessing proportionality, Federal Rule of Civil Procedure 26 directs consideration of “the importance of the issues at stake in the action, the amount in controversy, the parties’ relative access to relevant information, the parties’ resources, the importance of the discovery in resolving the issues, and whether the burden or expense of the proposed discovery outweighs its likely benefit.” Fed. R. Civ. P. 26(b)(1). Helpfully, Judge Chhabria provided further guidance at the March 5, 2020 Case Management Conference, stating:

I am concerned that Facebook has, you know, often made statements reflecting an unduly narrow view of what should be turned over to the Plaintiffs. And, you know, this is a big case. I mean, there is often a lot of talk about proportionality and whatnot. This is a big case. It is a significant issue. You know, and there is -- this is not the type of case where we are going to be saying: Well, that might end up -- that effort might end up uncovering some relevant information; but, you know, it is just too expensive or difficult, and so we are not going to make Facebook do it. This is really not one of those cases where that is very -- that type of argument is likely to carry the day. You know, and, as I have said a number of times, you know, the best way to figure out what happened as it relates to the claims that are going forward now is to -- for Facebook to produce all information, all documents about the practices associated with giving third parties access to friends' information and friends' of friends information.

Tr. at 28:25-29:18. Judge Chhabria’s observations regarding the size of this case remain on point. The proposed class period extends from 2007 to the present, the potential class members number in the hundreds of millions, and the third parties with whom Facebook shared user data appear to number in the tens of thousands. In that context, Plaintiffs’ request for the data concerning ten individual users seems not only proportional to the needs of the case but modest.

Furthermore, Facebook’s claims of burden are unsupported. “[T]he party opposing discovery has the burden of showing that discovery should not be allowed, and also has the burden of clarifying, explaining and supporting its objections with competent evidence.” *Harris v. Best Buy Stores, L.P.*, No. 3:15-cv-00657-HSG (KAW), 2016 WL 6024556, at \*1 (N.D. Cal. Oct. 14, 2016) (quoting *La. Pac. Corp. v. Money Mkt. 1 Institutional Inv. Dealer*, 285 F.R.D. 481, 485 (N.D. Cal. 2012)). A party claiming undue burden or expense “ordinarily has far better information—perhaps the only information—with respect to that part of the determination.” Fed. R. Civ. P. 26(b)(1) advisory committee’s note (2015). Therefore, the “party claiming that

discovery imposes an undue burden must ‘allege specific facts which indicate the nature and extent of the burden, usually by affidavit or other reliable evidence.’” *Sullivan v. Personalized Media Commc’ns, LLC*, No. 16-MC-80183-MEJ, 2016 WL 5109994, at \*3 (N.D. Cal. Sept. 21, 2016) (quoting *Nationstar Mortg., LLC v. Flamingo Trails No. 7 Landscape Maint. Ass’n*, No. 2:15-cv-01268-RFB-NJK, 2016 WL 4071988, at \*4 (D. Nev. July 28, 2016)).<sup>8</sup> Facebook has furnished no evidentiary support for its objections of undue burden and its objections should be overruled.

Plaintiffs emphasize that they are seeking discovery about *ten Named Plaintiffs*—not millions, not thousands, and not hundreds of users. Based on the information Plaintiffs obtain about themselves, and about Facebook’s general practices and procedures, they will seek to prove their class claims. Facebook’s contention that Plaintiffs are not even entitled to obtain in discovery the evidence necessary to show what Facebook collects about them, and with whom it shares the information is impossible to square with Facebook’s basic discovery obligations under the Federal Rules.

### III. CONCLUSION

For the reasons set forth above, Plaintiffs respectfully request that the Court deny Facebook’s motion to impose a discovery stay and grant Plaintiffs’ motion to compel discovery responsive to Requests for Production Nos. 9 through 13.

---

<sup>8</sup> See also *SPS Techs., LLC v. Briles Aerospace, Inc.*, No. CV 18-9536 MWF (ASx), 2020 WL 4341717, at \*2-3 (C.D. Cal. June 25, 2020) (overruling objection to requests for production of documents and noting that the party resisting discovery must describe “in specific detail, how each Request is overly broad and unduly burdensome by submitting affidavits or other evidence describing the nature of the burden”); *Polaris Innovations Ltd. v. Kingston Tech. Co.*, No. CV1600300CJCRAOX, 2017 WL 3275615, at \*6 (C.D. Cal. Feb. 14, 2017) (court grants motion to compel production of documents by defendant Kingston in part because “[r]egarding its assertion that the requests are overly burdensome, Kingston has not submitted any evidentiary declaration to support this objection.”).

Dated: September 28, 2020

Respectfully submitted,

KELLER ROHRBACK L.L.P.

BLEICHMAR FONTI & AULD LLP

By: /s/ Derek W. Loeser  
Derek W. Loeser

By: /s/ Lesley E. Weaver  
Lesley E. Weaver

Derek W. Loeser (admitted *pro hac vice*)  
Lynn Lincoln Sarko (admitted *pro hac vice*)  
Gretchen Freeman Cappio (admitted *pro hac vice*)  
Cari Campen Laufenberg (admitted *pro hac vice*)  
David Ko (admitted *pro hac vice*)  
Benjamin Gould (SBN 250630)  
Adele A. Daniel (admitted *pro hac vice*)  
1201 Third Avenue, Suite 3200  
Seattle, WA 98101  
Tel.: (206) 623-1900  
Fax: (206) 623-3384  
dloeser@kellerrohrback.com  
lsarko@kellerrohrback.com  
gcappio@kellerrohrback.com  
claufenberg@kellerrohrback.com  
dko@kellerrohrback.com  
bgould@kellerrohrback.com  
adaniel@kellerrohrback.com

Lesley E. Weaver (SBN 191305)  
Anne K. Davis (SBN 267909)  
Matthew P. Montgomery (SBN 180196)  
Angelica M. Ornelas (SBN 285929)  
Joshua D. Samra (SBN 313050)  
555 12th Street, Suite 1600  
Oakland, CA 94607  
Tel.: (415) 445-4003  
Fax: (415) 445-4020  
lweaver@bfalaw.com  
adavis@bfalaw.com  
mmontgomery@bfalaw.com  
aornelas@bfalaw.com  
jsamra@bfalaw.com

Christopher Springer (SBN 291180)  
801 Garden Street, Suite 301  
Santa Barbara, CA 93101  
Tel.: (805) 456-1496  
Fax: (805) 456-1497  
cspringer@kellerrohrback.com

Eric Fierro (admitted *pro hac vice*)  
3101 North Central Avenue, Suite 1400  
Phoenix, AZ 85012  
Tel: (602) 248-0088  
Fax: (602) 248-2822  
efierro@kellerrohrback.com

*Plaintiffs' Co-Lead Counsel*

**ATTESTATION PURSUANT TO CIVIL LOCAL RULE 5-1(i)(3)**

I, Derek W. Loeser, attest that concurrence in the filing of this document has been obtained from the other signatory. I declare under penalty of perjury that the foregoing is true and correct.

Executed this 28th day of September, 2020, at Seattle, Washington.

/s/ Derek W. Loeser

Derek W. Loeser

### **CERTIFICATE OF SERVICE**

I, Sarah Skaggs, hereby certify that on September 28, 2020, I electronically filed the foregoing with the Clerk of the United States District Court for the Northern District of California using the CM/ECF system, which shall send electronic notification to all counsel of record.

**In addition, the following were served via email:**

Paven Malhotra  
Matan Shacham  
Bryn Anderson Williams  
pmalhotra@keker.com  
bwilliams@kvn.com  
bwilliams@kvn.com

Anjeza Hassan  
annie.sara@yahoo.com

*/s/ Sarah Skaggs*

\_\_\_\_\_  
Sarah Skaggs

# Exhibit E

GIBSON, DUNN & CRUTCHER LLP  
Orin Snyder (*pro hac vice*)  
osnyder@gibsondunn.com  
200 Park Avenue  
New York, NY 10166-0193  
Telephone: 212.351.4000  
Facsimile: 212.351.4035

Kristin A. Linsley (SBN 154148)  
klinsley@gibsondunn.com  
Martie Kutscher (SBN 302650)  
mkutscherclark@gibsondunn.com  
555 Mission Street, Suite 3000  
San Francisco, CA 94105-0921  
Telephone: 415.393.8200  
Facsimile: 415.393.8306

*Attorneys for Defendant Facebook, Inc.*

GIBSON, DUNN & CRUTCHER LLP  
Deborah Stein (SBN 224570)  
dstein@gibsondunn.com  
333 South Grand Avenue  
Los Angeles, CA 90071-3197  
Telephone: 213.229.7000  
Facsimile: 213.229.7520

Joshua S. Lipshutz (SBN 242557)  
jlipshutz@gibsondunn.com  
1050 Connecticut Avenue, N.W.  
Washington, DC 20036-5306  
Telephone: 202.955.8500  
Facsimile: 202.467.0539

**UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA  
SAN FRANCISCO DIVISION**

IN RE: FACEBOOK, INC. CONSUMER  
PRIVACY USER PROFILE LITIGATION,

This document relates to:

ALL ACTIONS

CASE NO. 3:18-MD-02843-VC

**DEFENDANT FACEBOOK, INC.'S  
REPLY BRIEF IN SUPPORT OF ITS  
REQUEST TO ENFORCE THE PARTIAL  
STAY OF DISCOVERY IN PRETRIAL  
ORDER NO. 20**

Judge: Hons. Vince Chhabria and  
Jacqueline Scott Corley  
Courtroom 4, 17th Floor



## TABLE OF CONTENTS

	<u>Page</u>
INTRODUCTION .....	1
ARGUMENT .....	2
I.    The stay imposed by Pretrial Order 20 includes a discovery stay. ....	2
II.   This case is about information users share with their friends on Facebook.....	3
A.    The four live theories all concern data users shared with their Facebook friends.....	3
B.    The threshold “global issues” addressed in the Order show that the actionable claims relate only to information users shared on Facebook.....	4
III.  Facebook produced all data Plaintiffs shared on Facebook; no other user data is relevant. ....	6
IV.   The Court should deny Plaintiffs’ “Cross-Motion to Compel.” .....	10
CONCLUSION .....	10

## TABLE OF AUTHORITIES

	<u>Page(s)</u>
 <b>Cases</b>	
<i>Meyers v. Cty. of Sacramento</i> , 2020 WL 207213 (E.D. Cal. Jan. 14, 2020).....	2
<i>Mujica v. AirScan, Inc.</i> , 771 F.3d 580 (9th Cir. 2014).....	3
<i>WPP Luxembourg Gamma Three Sarl v. Spot Runner, Inc.</i> , 655 F.3d 1039 (9th Cir. 2011).....	2

## INTRODUCTION

The lawsuit Plaintiffs describe is not this case. This case is about information sharing. Specifically, it concerns sensitive information that users shared with their Facebook friends and that third parties allegedly accessed as a result of friend sharing, whitelisting, and integration partner agreements. Pretrial Order 20 is clear on this point, and Plaintiffs do not identify a single line in Judge Chhabria's comprehensive order, much less in their own allegations, that supports their description of the case that survived dismissal.

The Order explains on its first page: "This lawsuit . . . is about Facebook's practice of sharing its users' personal information with third parties." Dkt. 298 ("Order") at 1 (emphasis added). It then says that each of the four live theories concerns "substantive and revealing content that users intended only for a limited audience [*i.e.*, their Facebook friends], such as their photographs, videos they made, videos they watched, their religious and political views, their relationship information, and the actual words contained in their messages." *Id.*; see also *id.* at 7, 13, 17. The user data relevant to those theories consists of "information [users] make available to their friends on [Facebook]." *Id.* at 1.

Plaintiffs do not dispute that Facebook produced all of the information the Named Plaintiffs ever shared on Facebook. These productions consist of more than one million pages of data and necessarily include any data Facebook shared under the live theories. But, Plaintiffs insist they are entitled to any other data that has ever crossed Facebook's servers that relates in any way to any Named Plaintiff and all derivative materials drawing on this data. Plaintiffs seek these materials even if the underlying data is not associated with any user and even if they were never shared with any third party. Plaintiffs do not even attempt to explain why they would need such data in a case concerning information they shared on the Facebook platform and that Facebook allegedly shared beyond the audience Plaintiffs intended. Instead, Plaintiffs openly admit that they seek these extraneous materials not to pursue live claims, but to resuscitate stayed and dismissed theories or to search for new ones.

Plaintiffs largely avoid the Court's instruction to brief "what the scope of discovery is based on the claims in Judge Chhabria's ruling [Pretrial Order 20]." 9/4/2020 Hr'g Tr. at 5:8-10. Instead, Plaintiffs devote the majority of their brief to side issues and seek to compel Facebook to produce all documents responsive to five RFPs that are not before the Court. The Court should disregard these diversions, conform discovery to the four operative theories, and deny Plaintiffs' motion to compel.

## ARGUMENT

### I. The stay imposed by Pretrial Order 20 includes a discovery stay.

Plaintiffs take the surprising position that Pretrial Order 20 sets virtually no bounds on the scope of discovery in this case and allows them to explore theories Judge Chhabria stayed or dismissed.<sup>1</sup>

Plaintiffs’ position makes no sense. When a stay is in place, it “include[s] a stay of discovery.” *Meyers v. Cty. of Sacramento*, 2020 WL 207213, at \*1 (E.D. Cal. Jan. 14, 2020). Judge Chhabria stayed all but Plaintiffs’ core theories because Plaintiffs filed a 1,440-paragraph pleading. As he explained, “it seems the plaintiffs sought to identify anything Facebook has ever been reported to have done wrong . . . [and] the presence of so many disparate and vague allegations ma[de] it nearly impossible for Facebook to meaningfully respond to all of them, much less for the Court to effectively address them.” Order at 5-6. In order to avoid “boggling the case down at the pleading stage for years,” *id.* at 6, Judge Chhabria therefore issued an opinion regarding Plaintiffs’ core allegations, without addressing most of their improperly pleaded theories, which he stayed, *id.* Judge Chhabria surely did not intend to allow discovery on hundreds of “disparate and vague allegations” that did not satisfy Rule 8. The very point of the stay was to focus this case—not to allow Plaintiffs to explore “anything Facebook has ever been reported to have done wrong” without stating cognizable claims.

Plaintiffs even suggest Pretrial Order 20 allows discovery to “reviv[e]” claims dismissed with prejudice. Opp’n at 6 n.3. To support this curious position, Plaintiffs cite a footnote in Judge Chhabria’s analysis of Plaintiffs’ deceit by concealment claim. *Id.* (citing Order at 37 n.21). Judge Chhabria held that Plaintiffs stated a plausible claim arising from Facebook’s alleged practices concerning whitelisting and integration partners. But he held the claim did not satisfy Rule 9(b)’s heightened pleading standard with respect to friend sharing and Facebook’s enforcement measures. He then said in a footnote that dismissal of a subset of the claim would not “preclude . . . plaintiff[s] from seeking revival if discovery reveals a factual basis that justified reconsideration.” Order at 37 n.21. Judge Chhabria cited *WPP Luxembourg Gamma Three Sarl v. Spot Runner, Inc.*, 655 F.3d 1039, 1059 (9th Cir. 2011), which holds a plaintiff who fails to satisfy the PLRA’s heightened pleading standard may potentially seek revival if other case-related discovery later allows the plaintiff to satisfy the

---

<sup>1</sup> In addition to the discovery Plaintiffs seek from Facebook, Plaintiffs have served overbroad subpoenas on 27 third parties. These parties also require clear guidance as to the scope of discovery.

PLRA’s heightened pleading standard. Judge Chhabria certainly did not intend this footnote to create a gaping hole allowing discovery on hundreds of allegations that did not survive dismissal. *See Mujica v. AirScan Inc.*, 771 F.3d 580, 593 n.7 (9th Cir. 2014) (“To the extent [earlier cases] suggest[] that courts retain discretion to permit discovery whenever a plaintiff has failed to satisfy Rule 8’s plausibility standard, it is simply incompatible with *Iqbal* and *Twombly*.”).

Pretrial Order 20 plainly defines the scope of discovery in defining the scope of the case.

## **II. This case is about information users share with their friends on Facebook.**

A plain reading of Pretrial Order 20 explains the scope of the case Judge Chhabria allowed to move forward.<sup>2</sup> Plaintiffs say the Order describes this case as concerning any data Facebook receives or infers about users and how that data may be used to target them. To support this position Plaintiffs quote vague passages from the Order stating the case concerns “sensitive information.” Plaintiffs then say Judge Chhabria did not define “sensitive” and ask the Court to interpret the term to include any data Plaintiffs believe to be personal—including information they provide to third parties, information third parties collect through cookies, public records, and even inferences Facebook draws. Opp’n at 4.

Plaintiffs disregard what Pretrial Order 20 actually says. It describes “sensitive information” to be “substantive and revealing content that users intended only for a limited audience,” and clarifies that this data is “**information [users] make available to their friends on [Facebook].**” Order at 1. To read the ruling otherwise would expand the case far beyond what Judge Chhabria considered and would also raise a host of thorny legal questions his Order does not address.

### **A. The four live theories all concern data users shared with their Facebook friends.**

As discussed, Pretrial Order 20 allows four theories of relief to move forward. Each theory concerns information users shared with their Facebook friends.

**Friend sharing.** Friend sharing was a capability through which users could share with apps information their friends posted and made available to their Facebook friends. Plaintiffs do not dispute

<sup>2</sup> Plaintiffs disingenuously argue that Facebook takes an “unduly narrow” view of discovery, citing a comment Judge Chhabria made before discovery began advising Facebook to produce materials regarding “friends’ information and friends’ of friends information.” Opp’n at 13. Facebook has now produced nearly 1.5 million pages of documents, before the parties have even reached a search term agreement, including all information the Named Plaintiffs shared with their friends and friends of their friends. Those productions also include all of the Facebook documents produced to the FTC in response to its document requests in two related investigations. They also include documents produced to a host of other government actors in related actions responsive to Plaintiffs’ RFPs. In addition to these materials, Facebook proposed search terms hitting on millions of additional documents.

this. Friend sharing underlies the Cambridge Analytica events, it has been hotly litigated, and there is no dispute as to what it is about. The Order explains: “[W]hen users accessed apps on the Facebook Platform, the app developers were not merely able to obtain information about the users they were interacting with; [but] were also able to obtain any information about the users’ Facebook friends that the users themselves had access to,” “such as photographs, videos they watched, religious preferences, posts, and even sometimes private one-on-one messages sent through Facebook.” *Id.* at 6-7.

**Whitelisting.** Whitelisting is an extension of friend sharing and is about the same data. *Id.* at 8.

**Integration Partner Agreements.** Facebook allegedly “[s]har[ed] sensitive user information with business partners,” through a list of “integration partnerships,” to integrate Facebook with devices, websites, and social-media platforms. *Id.* at 8-9. As with the other theories, the “sensitive user information” at issue is “substantive and revealing content that users intended only for a limited audience [i.e. their Facebook friends].” *Id.* at 1. The purpose of these agreements was to allow users to integrate their Facebook activities *that they shared* on Facebook with other platforms and sites.

Plaintiffs say the Order allows claims relating to integration partners to proceed as to some broader set of “sensitive information” that they find personal in nature. Opp’n at 7. Plaintiffs provide no support for this assertion; the Order describes this theory as involving the same “sensitive user data” underlying the other theories of relief. And it must. As discussed below, the Order holds that Plaintiffs demonstrated standing, a reasonable expectation of privacy, and a lack of consent only with respect to Facebook’s alleged practice of sharing information users shared with their Facebook friends.

**Enforcement.** This theory relates to how Facebook enforced its data-use policies with respect to data third parties obtained through friend sharing, whitelisting, and integration partner agreements, and it concerns the same data that users shared with their Facebook friends. Order at 9.

**B. The threshold “global issues” addressed in the Order show that the actionable claims relate only to information users shared on Facebook.**

Pretrial Order 20 addressed various “global issues” and holds Plaintiffs demonstrated a reasonable expectation of privacy, standing, and lack of consent only with respect to Facebook’s alleged practice of sharing with third parties information users shared with their friends on Facebook.

**Expectation of privacy.** Pretrial Order 20 addresses Facebook’s argument that Plaintiffs were not injured, and therefore lack standing, because they did not have a reasonable expectation of privacy

over information they share with their Facebook friends. *Id.* at 1 (“Facebook argues that people have no legitimate privacy interest in information they make available to their friends on social media.”). With respect to users’ privacy expectations, Pretrial Order 20 holds: “the issue of whether users have a reasonable expectation of privacy *in information they share with their social media friends* is best understood as relating to the merits, not standing.” *Id.* at 10-11 n.2 (emphasis added). On the merits, the Order holds that “[w]hen you share sensitive information with a limited audience . . . you retain privacy rights and can sue someone for violating them.” *Id.* at 2. It then analyzes whether users retain a reasonable expectation of privacy over information *they share with their friends*, *see id.* at 10-12, and concludes: “social media users can have their privacy invaded if sensitive information *meant only for a few dozen friends* is shared more widely,” *id.* at 11.

Pretrial Order 20 is so clear that this case concerns information that users shared with their Facebook friends that it goes out of its way to say *sua sponte*: “It seems quite possible that a user whose settings allow information to be shared not only with their friends, but friends of friends, loses any expectation of privacy.” *Id.* at 11 n.3. Nowhere does Pretrial Order 20 consider whether users maintain a reasonable expectation of privacy over information beyond what users share on Facebook (as Plaintiffs wrongly suggest) such as information users provide third parties, public records, information third parties obtain through cookies, or information Facebook “infers” about users.

**Standing.** With respect to standing, Pretrial Order 20 holds: “The alleged injury is ‘concrete’ largely for the reasons already discussed – if you use a company’s social media platform *to share sensitive information with only your friends*, then you suffer a concrete injury when the company disseminates that information widely.” *Id.* at 17. The Order goes on to say that Plaintiffs’ injuries are sufficiently particularized with respect to which third parties allegedly received their data because, “[i]f, as alleged in the complaint, *Facebook made users’ ‘friends only’ information readily available* to such a broad swath of companies . . . it is virtually inevitable that some of these companies obtained information on the named plaintiffs.” *Id.* at 18. The Order did not hold—or even consider—whether Plaintiffs have standing to sue Facebook with respect to information users did not share on Facebook.

**Consent.** On the issue of consent, the Court addressed whether Plaintiffs consented to the conduct underlying their claims because they “agreed, when they signed up for their accounts, that



Facebook could disseminate their ‘friends only’ information in the way it has done.” *Id.* at 18. Pretrial Order 20 holds that judicially noticeable materials demonstrate that a subset of users consented to sharing their “friends only” information through friend sharing, but do not establish at the pleading stage that all users consented to sharing friends-only information through friend sharing, whitelisting, and integration agreements. *Id.* at 18-29. The Order did not consider whether Plaintiffs consented to sharing information they did not share on Facebook.

The Order is clear that this case is about sensitive information users made available to their friends on Facebook and third parties allegedly accessed. Discovery must conform to these theories.

### **III. Facebook produced all data Plaintiffs shared on Facebook; no other user data is relevant.**

Facebook produced *more than one million pages of content and information* related to the Named Plaintiffs.<sup>3</sup> *Those materials include everything each Named Plaintiff ever shared on Facebook* (unless they deleted it). This includes, but is not limited to, the “Download Your Information” (“DYI”) file that Facebook makes available to users,<sup>4</sup> plus *additional* information.

Plaintiffs do not dispute that the produced materials include any data users shared with their Facebook friends (sensitive or otherwise). Yet, Plaintiffs demand that Facebook search *millions* of disaggregated data sets for any data to have ever crossed Facebook’s systems relating to a Named Plaintiff and any derivative materials drawing on that data—such as data sets tracking hours of peak user activity to monitor strains on Facebook’s system. They demand such materials without regard for whether they were shared with any third party, much less under a live theory. To support this position, Plaintiffs misinterpret a handful of Facebook documents,<sup>5</sup> but their argument boils down to the following: Facebook has documents drawing on data relating to users; therefore, Facebook must search

<sup>3</sup> Since filing its opening brief, Facebook produced approximately 250,000 additional pages of information related to Named Plaintiffs who were added to the case in August.

<sup>4</sup> Plaintiffs assert that the DYI data is not useful because it does not display on an item-by-item basis the audience that Plaintiffs set for each of their posts. Facebook agreed to investigate whether it could produce this data for relevant posts—bearing in mind that the request involves granular data for *more than a million pages* of activity. Facebook also reminded Plaintiffs that their accounts display this information. If Plaintiffs believe the audience set to a particular post is critical evidence for their case, they could screen-shot that information from their accounts and produce it. They could also identify particular posts to Facebook so that Facebook can produce the relevant information.

<sup>5</sup> Because the Court ordered the parties not to submit declarations or evidence, 9/4/2020 Hr’g Tr. at 5:8-10, 18-22, Facebook does not here submit declarations or documents to dispute Plaintiffs’ characterization of the materials they cite. If the Court is inclined to issue a ruling relying on the exhibits Plaintiffs submitted, Facebook respectfully requests permission to do so.



for and produce any materials drawing upon any data it has ever collected or created that relates in any way to a Named Plaintiff. The Court should reject this position, which largely asks Facebook to search for materials that are out of scope and consist largely of data already produced in other formats.

**Data the Named Plaintiffs did not share on Facebook is out of scope**, including public records, data Plaintiffs shared with third parties, and information created by Facebook. Facebook produced all of the information the Named Plaintiffs shared on Facebook (sensitive or otherwise). These productions necessarily include any information shared under the live theories.<sup>6</sup>

Plaintiffs say the case is about data Facebook creates and that third parties share with Facebook that is used to draw “inferences” about users. For instance, Plaintiffs may allow websites to collect data about their shopping habits through cookies. Those sites then might share this data with other parties (including Facebook) to better place the site’s advertisements. As discussed above, nothing in Pretrial Order 20 supports Plaintiffs’ argument that this type of data is part of this lawsuit. This would be a very different case if—as Plaintiffs say—it were about Facebook sharing information that third parties passed on to Facebook. To establish this sort of “third party data” claim, Plaintiffs would have had to allege (and prove) the nature of each of their relationships with the specific third parties at issue, the circumstances under which those third parties obtained their data, whether each individual user consented to that third party sharing data with Facebook, the circumstances under which the data was provided to Facebook, and so on. None of that is at issue here and nothing in Pretrial Order 20 suggests it is. Nor could Plaintiffs conceivably establish facts of this nature on a class-wide basis.

Plaintiffs seem to concede they demand these materials because “the very purpose of collecting all of this data . . . is to use it to target users.” Opp’n at 12. As Plaintiffs admit, Pretrial Order 20 dismisses their targeted advertising theory.<sup>7</sup> To the extent any advertisers received sensitive user data through friend sharing, whitelisting, or integration agreements, that user data was already produced.

<sup>6</sup> To describe the data Plaintiffs believe Facebook maintains, Plaintiffs cite their Exhibit B at page 9, which was prepared by an employee in 2014 and regards Facebook’s ads platform. The document does not reflect Facebook’s standard terminology, nor does Facebook agree with Plaintiffs’ characterization of the document. In any case, Facebook does not dispute that it receives data from third parties in connection with its ads platform and maintains internal analyses which rely on user data.

<sup>7</sup> Plaintiffs walk back their position that they need discovery to pursue their dismissed “targeted advertising” and “psychographic marketing” theories. Opp’n at 11. But Plaintiffs have been arguing for a year that these theories justify their demands for every piece of information Facebook collects and infers about users and took this position in the recent joint status updates that prompted this briefing. See 8/13/2020 Status Update at 2-3, 9, Dkt. 495; 7/30/2020 Status Update at 3, Dkt. 484.

In any case, as discussed in its opening brief and below, Facebook actually produced the majority of data it receives from third parties in the off-Facebook Activity portion of the DYI materials.

**Data not shared through one of the four theories is out of scope.** Plaintiffs say they provide evidence that Facebook shares data beyond what users share on Facebook. Opp’n at 9-10. Even if that were true, it is not relevant to this lawsuit. The live theories concern data users shared with their Facebook friends that third parties accessed via friend sharing, whitelisting, or integration agreements.

In any case, the documents Plaintiffs cite describe Facebook’s data *sources*; they say nothing about whether or how Facebook *shares* information. Of note, Plaintiffs claim their Exhibit C “confirms that Facebook shares [the data they seek] with third parties.” Opp’n at 9. Exhibit C is an email outlining *hypothetical* platform capabilities—it does not discuss what data Facebook actually shared.

**The integration partner theory does not entitle Plaintiffs to all data from third parties.** Plaintiffs suggest Facebook must locate and produce all data points it has ever received from any third party regarding a Named Plaintiff because Facebook’s integration partner agreements were built in part on “data reciprocity.” Opp’n at 11. This argument is a red herring and misrepresents what “data reciprocity” means. Facebook did not, as Plaintiffs suggest, have agreements with integration partners to trade user data. Data reciprocity arrangements allowed users to post their Facebook activities to third-party platforms if the third-party platform also allowed its users to post their activities to Facebook. Plaintiffs acknowledge this. *See* SACC ¶ 657(g) (“‘Reciprocity’ agreements . . . requir[ed] Apps that used data from Facebook to allow their users to share their data back to Facebook”); *see also id.* ¶¶ 239, 745. Any user data relating to that type of sharing was produced. Again, Facebook produced everything the Plaintiffs shared on Facebook. This includes any Facebook activities Plaintiffs *elected* to share on other platforms and any off-Facebook activities Plaintiffs *elected* to share on Facebook. In any event, even if some other data from integration partners existed, only data received from *those partners* could even possibly be relevant—not data from thousands of other third parties.<sup>8</sup>

**Plaintiffs’ SCA and VPPA claims do not require additional data.** Plaintiffs contend this case concerns data beyond what they shared on Facebook because Pretrial Order 20 did not dismiss

---

<sup>8</sup> Plaintiffs concede they seek any such data to prove damages. If the Court is inclined to require broad discovery to support damages, Facebook respectfully requests the opportunity to submit briefing regarding why any such discovery should be bifurcated from liability-related discovery.

their claims under the Stored Communications Act (“SCA”) and Video Privacy Protection Act (“VPPA”). Opp’n at 4-5. Plaintiffs highlight that their SCA claim turns, in part, on “whether the shared information includes the contents of an electronic communication.” *Id.* at 5. But the sensitive data at issue includes “private one-on-one messages” sent on Facebook. Order at 17; *see id.* at 1, 32. Similarly, Plaintiffs’ VPPA claim survived dismissal on the basis that Facebook shared “information about the videos that users received in their private [Facebook] messages and about videos they ‘liked.’” *Id.* at 34. Plaintiffs’ messages and any videos they shared or liked were produced.

**The additional data Plaintiffs seek cannot even be reasonably collected.**<sup>9</sup> Facebook understands Plaintiffs seek two forms of data: (i) any additional data regarding users’ off-Facebook Activity provided by third parties, and (ii) any derivative materials that draw from user data. Again, these materials are not relevant to any live theory. Facebook also cannot reasonably identify them.

With respect to off-Facebook Activity, as Facebook explained in its opening brief, the produced DYI materials include the vast majority of data Facebook receives from third parties. It is not clear what else Plaintiffs seek or how it could be relevant. Any off-Facebook Activity provided by third parties that is *not* included in the DYI materials is data Facebook has not linked to a particular user or data that is so granular that it is preserved only temporarily. There is no centralized way to search for either type of data. To the extent it exists, it is organized by the third parties who provided it. Facebook would therefore need to review every data set it has received from thousands of third parties and then attempt to link to the Named Plaintiffs data points it previously did not associate with any user. Such an exercise is unlikely to be fruitful or at all useful, particularly on a class-wide basis.

Facebook also explained that, within 90 days, any user data not included in the DYI materials is disassociated from the user’s ID, anonymized entirely, or deleted (depending on the nature of the data and any business reasons for retaining it). Plaintiffs argue that Facebook should *still* be able to find any derivative materials drawing from data relating to any Named Plaintiff because data disassociated from a user’s ID can sometimes be linked back to the user’s account. Plaintiffs’ explanation of this process is oversimplified, incorrect, and ignores that much of the data they demand

---

<sup>9</sup> Plaintiffs say Facebook did not prove undue burden because it did not submit declarations or evidence. The Court instructed the parties not to submit such materials. 9/4/2020 Tr. at 5:8-10, 18-22.

is fully anonymized or not retained at all.<sup>10</sup> The argument also misses the point. There is no way for Facebook to run a centralized search for a user's ID, random ID, or any "hashed data" identifiers across millions of data sets, which are largely used for business analytics (like scoping infrastructure needs). The only way to search these tables is to open all of them and search each to find any data relating to a particular user, whether by user ID or otherwise. The issue is opening and searching each table.<sup>11</sup>

To be clear, Facebook is not—as Plaintiffs suggest—urging the Court to issue a ruling regarding the scope of discovery based on undue burden. Facebook is highlighting that this is not a situation in which there are marginally relevant materials that are easy to sweep into an ongoing collection. The user data Plaintiffs seek has nothing to do with the four operative theories, most of it was actually produced, and any additional data would be virtually impossible to locate. If Plaintiffs are able to identify some specific type of data about user activity that is relevant to the case, Facebook will search for that data. But Plaintiffs' position that Facebook must search the entire company for every document including any data relating to a Named Plaintiff is simply not reasonable.

#### **IV. The Court should deny Plaintiffs' "Cross-Motion to Compel."**

Plaintiffs style their brief as a "cross-motion to compel" compliance with five RFPs and criticize Facebook for not submitting declarations and evidence about these requests. The Court should disregard this diversion, which puts the cart before the horse. The Court directed the parties to submit "no declarations," as this briefing is "just a legal question as to what the scope of discovery is based on the claims in Judge Chhabria's ruling." 9/4/2020 Tr. at 5:8-10, 18-22. Facebook told Plaintiffs it will produce materials responsive to the RFPs they identify that are in Facebook's possession and relate to the operative theories. The Court must resolve this threshold legal issue before it can consider (on a full evidentiary record) whether Facebook produced the relevant evidence responsive to specific RFPs.

### **CONCLUSION**

The Court should enforce the stay Judge Chhabria imposed, allow discovery only on the four operative theories of relief detailed in Pretrial Order 20, and deny Plaintiffs' cross-motion to compel.

<sup>10</sup> Plaintiffs' Exhibit B, which they cite on page 12 of their brief, describes the ability to reidentify data points that remain live on a user's Facebook page. This live data has already been produced.

<sup>11</sup> Plaintiffs did not ease the burden of searching millions of data sets by identifying 10 Named Plaintiffs they *intend* to identify as class representatives. In any case, the other 14 Named Plaintiffs have not withdrawn their claims and have reserved their rights to proceed as class representatives.

DATE: October 8, 2020

Respectfully submitted,

**GIBSON, DUNN & CRUTCHER, LLP**

By: Orin Snyder

Orin Snyder (*pro hac vice*)

osnyder@gibsondunn.com

GIBSON, DUNN & CRUTCHER LLP

200 Park Avenue

New York, NY 10166-0193

Telephone: 212.351.4000

Facsimile: 212.351.4035

Deborah Stein (SBN 224570)

dstein@gibsondunn.com

GIBSON, DUNN & CRUTCHER LLP

333 South Grand Avenue

Los Angeles, CA 90071-3197

Telephone: 213.229.7000

Facsimile: 213.229.7520

Joshua S. Lipshutz (SBN 242557)

jlipshutz@gibsondunn.com

GIBSON, DUNN & CRUTCHER LLP

1050 Connecticut Avenue, N.W.

Washington, DC 20036-5306

Telephone: 202.955.8500

Facsimile: 202.467.0539

Kristin A. Linsley (SBN 154148)

klinsley@gibsondunn.com

Martie Kutscher (SBN 302650)

mkutscherclark@gibsondunn.com

GIBSON, DUNN & CRUTCHER LLP

555 Mission Street, Suite 3000

San Francisco, CA 94105-0921

Telephone: 415.393.8200

Facsimile: 415.393.8306

*Attorneys for Defendant Facebook, Inc.*

# Exhibit F

UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA

IN RE: FACEBOOK, INC. CONSUMER  
PRIVACY USER PROFILE LITIGATION

MDL No. 2843  
Case No. 18-md-02843-VC (JSC)

**DISCOVERY ORDER NO. 9**  
(Dkt. Nos. 515, 526, 537, 548)

This MDL matter has been assigned to the undersigned for management of discovery. Now pending before the Court are the Parties' briefs concerning the proper scope of discovery related to the data Facebook accumulates about the named Plaintiffs. (Dkt. Nos. 515, 526, 537, 548.) In brief, Facebook contends that the district court's order specifically defined the data at issue as "substantive and revealing content that users intended only for a limited audience." (Dkt. No. 298.) Based on this definition, Facebook argues that for any named Plaintiff data to be relevant and discoverable, it must meet two criteria. First, the discoverable data must have arisen from user activity occurring on the Facebook platform, such as Facebook posts and sent messages. Second, the named Plaintiff must have then overtly shared such data with a limited audience, such as their friends. Facebook submits that this is the only plausible reading of the district court's order limiting Plaintiffs to four actionable categories of potential liability. Plaintiffs respond that the universe of discoverable data Facebook collects for each user is much larger and necessarily includes: (1) user activity occurring off the Facebook platform; and (2) user data that can be inferred from user activity occurring on or off the Facebook platform. A second question presented by the briefs is whether discovery may proceed on the claims the district court stayed.

After carefully considering the papers submitted by the Parties, and consulting with the district court, the Court rules that discovery is not as limited as Facebook contends. Plaintiffs correctly argue that Facebook's restrictive view of relevant discovery would exclude an enormous

1 amount of information that Facebook collects and shares with third parties about Facebook's  
2 users. The district court's order (Dkt. No. 298) did not limit Plaintiffs' claims to only challenging  
3 the sharing of data Facebook collects from a user's on-platform activity; the claims also challenge  
4 Facebook's sharing of user data and alleged failure to monitor how third parties used such shared  
5 information.

6 Accordingly, the Court rules the discoverable user data at issue includes:

- 7
- Data collected from a user's on-platform activity;
  - 8 • Data obtained from third parties regarding a user's off-platform activities; and
  - 9 • Data inferred from a user's on or off-platform activity.

10 As for the stayed claims, and again after consulting with the district court, the Court rules  
11 that discovery is stayed as to the stayed claims. Of course, if a particular discovery request is  
12 relevant to both a stayed and non-stayed claim, then discovery is not stayed merely because the  
13 discovery request is also relevant to a stayed claim.

14 **IT IS SO ORDERED.**

15 Dated: October 29, 2020

16  
17  
18   
JACQUELINE SCOTT CORLEY  
United States Magistrate Judge



# Exhibit G

**UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA  
SAN FRANCISCO DIVISION**

IN RE: FACEBOOK, INC. CONSUMER  
PRIVACY USER PROFILE LITIGATION,

This document relates to:

ALL ACTIONS

**MDL NO. 2843**

CASE NO. 3:18-MD-02843-VC-JSC

HON. VINCE CHHABRIA  
COURTROOM 4 – 17<sup>TH</sup> FLOOR  
SPECIAL MASTER, DANIEL GARRIE, ESQ.

**ORDER RE: PLAINTIFFS' MOTION TO  
COMPEL PRODUCTION OF PLAINTIFF  
DATA**

## INTRODUCTION

1. Pending before Special Master Garrie is Plaintiffs’ Motion to Compel Production of Named Plaintiffs’ Content and Information.

## BACKGROUND

2. On November 25, 2019, Plaintiffs’ served Requests for Production Nos. 9-13, which seek documents relating to the named Plaintiffs in this matter (“Named Plaintiffs”).<sup>1</sup> See Exhibit A. In brief, Request No. 9 seeks all documents relating to each of the Named Plaintiffs; Request No. 10 seeks documents sufficient to show the categories of content and information Facebook collects, tracks, and maintains about them; and Requests Nos. 11-13 seek documents identifying third parties that were able to access information about the Named Plaintiffs. Id.

3. In response to Requests for Production Nos. 9-13, Facebook produced more than one million pages of individual user data it maintained relating to the Named Plaintiffs, most of which was obtained from the “Download Your Information” tool (“DYI Tool”).<sup>2</sup> The data obtained from the DYI Tool is mostly limited to information pertaining to users’ on platform Facebook activity. See Exhibit B (List of DYI Tool Data Fields).

4. Statements by Facebook’s counsel during an August 14, 2020 discovery hearing indicated that Facebook maintained additional data related to the Named Plaintiffs that was not produced. See Exhibit C (8/14/2020 Discovery Hearing Transcript) at 8:10-13 (“There is other – there’s Facebook-generated information, information generated by third parties, information

---

<sup>1</sup> There were originally 30 named Plaintiffs, but this has been reduced to nine named Plaintiffs.

<sup>2</sup> The DYI Tool is a tool by which Facebook users can download certain pieces of information related to the user’s Facebook activity and related data. A list of the types of information that can be downloaded via the DYI Tool is provided in Exhibit B.

received from third parties. We have not represented that that is comprehensively included in our production.”).

5. Plaintiffs filed a motion last September to compel additional discovery related to Requests for Production Nos. 9-13. See Exhibit D (9/28/2020 Motion to Compel). Plaintiffs asked the Court to compel production of sensitive information Facebook derives and collects from business partners, app developers, apps, and other sources. This request included “native, appended and behavioral data” and purportedly anonymized data that could be connected to the Named Plaintiffs. Id. at 7-11.

6. On October 8, 2020, Facebook responded to Plaintiffs’ motion to compel. See Exhibit E (Facebook Opposition to Plaintiffs’ 9/28/2020 Motion to Compel). Facebook contended that all information related to the Named Plaintiffs that they did not themselves share on Facebook was outside the scope of the case; that all information not shared through one of the four theories of the case was not within the scope of the case; that Plaintiffs were not entitled to all data collected from third parties about the Named Plaintiffs; that the Stored Communications Act and Video Protection Privacy Act claims did not require the production of additional data Facebook had collected about the Named Plaintiffs; and that Facebook could not reasonably collect any of the additional information Plaintiffs sought. Id. at 6-10.

7. On October 29, 2020, Judge Corley issued Discovery Order No. 9, ruling “that discovery is not as limited as Facebook contends” and “the discoverable user data at issue includes: [1] Data collected from a user’s on-platform activity; [2] Data obtained from third parties regarding a user’s off-platform activities; and [3] Data inferred from a user’s on or off-platform activity.” See Exhibit F (Discovery Order No. 9) at 2.

8. In Discovery Order No. 11, Judge Corley provided further clarification on the discoverable user data intended to be included under Discovery Order No. 9:

It also contended that Plaintiffs conceded that user data not shared with or accessible to third parties is not relevant, (Dkt.No. 548 at 10), and because Facebook does not share inferred user data, the inferred user data Facebook maintains is not relevant. Facebook both collects and uses data about its users as part of its business model, including data derived from third parties. How it specifically uses this data is an open question, but if the Court were to accept Facebook's arguments about the scope of production, it would eliminate Discovery Order No. 9's third category of discovery: data inferred from a user's on or off-platform activity. What is needed now is more detail about Facebook's collection and use of user data so future discovery requests can be tailored to Plaintiffs' better understanding of the internal operations of Facebook as well the terminology it uses for describing data that is potentially responsive to Plaintiffs' discovery requests. See Exhibit G (Discovery Order No. 11) at 1.

9. Following Judge Corley's orders, Facebook did not produce additional documents in response to Requests for Production Nos. 9-13.

10. On October 6, 2021, Special Master Garrie and Judge Andler declared impasse on the issue of whether Facebook should be compelled to produce additional documents related to the Named Plaintiffs pursuant to Discovery Order No. 9.

11. On October 18, 2021, Plaintiffs submitted their opening brief to Special Master Garrie on this issue. See Exhibit H. Plaintiffs argue that (a) the court has already determined the information Plaintiffs seek is relevant—whether or not Facebook claims that it has been shared; (b) whether the Named Plaintiffs' information was shared is a contested question on which Plaintiffs are entitled to evidence; (c) Facebook has failed to substantiate a disproportionate burden in identifying the data it possesses relating to nine people; and (d) Plaintiffs have made proposals to reduce the burden of production on Facebook. Id.

12. On October 28, 2021, Facebook submitted its Opposition to Plaintiffs' Motion to Compel Production of Named Plaintiffs' Content and Information. See Exhibit I. Facebook argues,

among other things, that (a) the scope of discovery is limited to information Facebook shared with third parties; (b) Plaintiffs are judicially estopped from seeking information that was not shared; and (c) the information Plaintiffs now seek is nonresponsive and otherwise unavailable. Id.

13. On November 2, 2021, Plaintiffs submitted their Reply in which they argue, among other things, (a) Judge Corley's orders entitle Plaintiffs to the discovery they seek; (b) Plaintiffs are entitled to probe Facebook's assertion that it has already produced all the content and information it has shared or made accessible to third parties; (c) Plaintiffs are entitled to answers to Interrogatories 16 and 17; and (d) the relief Plaintiffs are requesting is intended to lighten Facebook's burden. See Exhibit J.

14. Facebook subsequently objected to Plaintiffs reply claiming that Plaintiffs introduced new arguments and evidence for the first time, in violation of the Discovery Protocol. See Exhibit K (Facebook's Response to Plaintiffs' Objection Regarding Named Plaintiffs' Data Briefing) ("Plaintiffs sought **new relief** and introduced **twelve new documents** that Plaintiffs suddenly claim show gaps in Facebook's productions.").

### FINDINGS

15. Special Master Garrie finds that Discovery Order No. 9 does not limit the scope of discoverable data related to the Named Plaintiffs to data that was shared with third parties, as Facebook contends, because Judge Corley's ruling contains no language indicating such a limitation: "Accordingly, the court rules the discoverable user data at issue includes: [1] Data collected from a user's on-platform activity; [2] Data obtained from third parties regarding a user's off-platform activities; and [3] Data inferred from a user's on or off-platform activity." See Exhibit F at 2.

16. Moreover, Judge Corley clarified that Facebook’s interpretation of Discovery Order No. 9 is not what Judge Corley intended: “How [Facebook] specifically uses this data is an open question, but if the Court were to accept Facebook’s arguments about the scope of production, it would eliminate Discovery Order No. 9’s third category of discovery: data inferred from a user’s on or off-platform activity.” See Exhibit G at 1.

17. Special Master Garrie finds that Facebook appears to maintain data related to the Named Plaintiffs that was not produced in response Requests for Production Nos. 9-13. See Exhibit C at 8:10-13 (“There is other – there’s Facebook-generated information, information generated by third parties, information received from third parties. We have not represented that that is comprehensively included in our production.”). For example, documents produced by Facebook indicate that Facebook collects data referred to as “Appended Data,” including public records, auto registration data, retail purchases, and credit card purchases, all of which fall into the second category of data from Discovery Order No. 9. See Exhibit L (FB-CA-MDL-00213424). However, Facebook has not produced this data as it is not available via the DYI Tool. See Exhibit B.<sup>3</sup>

18. Special Master Garrie finds that Plaintiffs requested new relief (answers to Interrogatories 16-17) and introduced new evidence (exhibits C, D, E, F, H, I, and J to Plaintiffs’ Reply) in their Reply brief in violation of the Discovery Protocol. Accordingly, Special Master Garrie did not consider this request for new relief or the new evidence items in reaching the findings herein.

//

---

<sup>3</sup> Facebook also appears to maintain data relating to the Named Plaintiffs’ on-platform activity that has not been provided, such as inferred interest and behavior data. See Exhibit L.

**ORDER**

19. No later than December 3, 2021, Facebook is to provide a list of data sources that may contain information related to the Named Plaintiffs pursuant to Discovery Order No. 9. The list of data sources is to include: (1) the name of the database or data log; (2) a description of the data source's purpose and function; and (3) a description of the types of Named Plaintiff data contained in the data source.

20. No later than December 10, 2021, the parties are to meet and confer and each submit to Special Master Garrie a proposed protocol for the production of Named Plaintiffs' data from the data sources identified by Facebook.

**IT IS SO ORDERED.**

Monday, November 29, 2021



\_\_\_\_\_  
Daniel Garrie  
Discovery Special Master



# Exhibit H


**Exhibit A to Pope Declaration**

**Data Systems That May Store or Interact With User Data**

[REDACTED]

[REDACTED]

[REDACTED]



# Exhibit I

HIGHLY CONFIDENTIAL - ATTORNEYS EYES ONLY

UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA

IN RE: ) MDL No. 2843  
)  
FACEBOOK, INC. CONSUMER ) Case No. 18-md-02843-  
PRIVACY USER PROFILE ) VC-JSC  
LITIGATION )  
----- )

HIGHLY CONFIDENTIAL

ATTORNEYS' EYES ONLY

SPECIAL MASTER: DANIEL GARRIE

REMOTE HEARING  
(Via Zoom Videoconference)  
Friday, January 14, 2022

REPORTED BY: Michelle Milan Fulmer  
CSR No. 6942, RPR, CRR, CRC

UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA

IN RE: ) MDL No. 2843  
)  
FACEBOOK, INC. CONSUMER ) Case No. 18-md-02843-  
PRIVACY USER PROFILE ) VC-JSC  
LITIGATION )  
----- )

Remote hearing taken before Michelle Milan  
Fulmer, a Certified Shorthand Reporter for the  
State of California, commencing at 12:08 p.m.,  
Pacific Standard Time, Friday, January 14, 2022.



1 APPEARANCES OF COUNSEL:

2  
3 FOR PLAINTIFFS:

4 BLEICHMAR FONTI & AULD  
5 BY: Lesley Weaver, Esq.  
Matthew Melamed, Esq.  
Anne K. Davis, Esq.  
6 555 12th Street, Suite 1600  
Oakland, California 94607  
7 TEL: (415) 445-4003  
EMAIL: lweaver@bfalaw.com  
8 mmelamed@bfalaw.com  
adavis@bfalaw.com  
9 (Via Zoom Videoconference)  
10

KELLER ROHRBACK, LLP  
11 BY: Derek W. Loeser, Esq.  
Cari Campen Laufenberg, Esq.  
12 1201 Third Avenue, Suite 3200  
Seattle, Washington 98101  
13 TEL: (206) 623-1900  
EMAIL: dloeser@kellerrohrback.com  
14 claufenberg@kellerrohrback.com  
(Via Zoom Videoconference)

15  
16 FOR DEFENDANT:

17 GIBSON DUNN & CRUTCHER  
BY: Martie Kutscher Clark, Esq.  
18 1881 Page Mill Road  
Palo Alto, California 94304-1211  
19 TEL: (650) 849-5348  
EMAIL: mkutscherclark@gibsondunn.com  
20 (Via Zoom Videoconference)  
21

GIBSON DUNN & CRUTCHER  
22 BY: Laura C. Mumm, Esq.  
200 Park Avenue  
23 New York, New York 10166-0193  
TEL: (212) 351-2404  
24 EMAIL: lmumm@gibsondunn.com  
(Via Zoom Videoconference)  
25

1 FOR DEFENDANT:

2 FACEBOOK LEGAL

BY: Ian Chen, Esq.

3 1 Hacker Way

Menlo, California 94025

4 EMAIL: ianchen0@fb.com

(Via Zoom Videoconference)

5  
6 ALSO PRESENT:

7 Michael Mann

(Via Zoom Videoconference)

HIGHLY CONFIDENTIAL - ATTORNEYS EYES ONLY

## I N D E X

WITNESS

EXAMINATION

DAVID POPE

PAGE

BY MR. GARRIE

6

## E X H I B I T S

Page

Exhibit

Description

Identified

Exhibit A

Spreadsheet

71

Page 5

1 Friday, January 14, 2022

2 12:08 p.m. - 3:21 p.m.

3 \*\*\*

4 MR. GARRIE: We're going on the record and  
5 we'll get started.

6 Court reporter, Ms. Milan Fulmer.

7

8 DAVID POPE,

9 called as a witness, having been first duly sworn by  
10 the Certified Shorthand Reporter, was examined and  
11 testified as follows:

12

13 EXAMINATION

14 BY MR. GARRIE:

15 Q Okay. Good afternoon, good morning,  
16 wherever you are in the world, Mr. Pope. Thank you  
17 for taking time out of your busy day to answer some  
18 questions.

19 I have a declaration that was submitted in  
20 this case regarding -- I'm sure you're aware of it,  
21 but just to make sure.

22 Did you submit a declaration on  
23 December 10th of 2021 in connection with the  
24 Facebook -- In Re: Facebook dispute?

25 A Yes.

Page 6

1 Q Okay. And maybe counsel for Facebook, or I  
2 can share my screen, but do you have that  
3 declaration in front of you? If not, I can share it  
4 just to make sure.

5 A Yeah. I do have the declaration on my  
6 computer.

7 Q Okay. All right. Sounds good.  
8 If you can just look at the declaration.  
9 Or here, I'll share my screen so everybody can see  
10 what I'm looking at.

11 So if you see here the declaration, in this  
12 declaration it says "My team," in Paragraph 3 right  
13 here, to inventory all data systems within Facebook  
14 and understand which of them retain user data and  
15 you spent a year working.

16 Do you see where I am?

17 A Yes. I see that.

18 Q Okay. And then you provided a list  
19 attached as Exhibit A.

20 A Uh-huh.

21 Q Okay. All you see right now is the pdf;  
22 right?

23 A That's correct.

24 Q Okay. So I have a high-level question  
25 before we --

1 Over the course of the year, right, you  
2 collect -- what exactly -- because I don't want to  
3 ask you for things that you don't have or that I  
4 don't understand.

5 So if you could just tell me sort of at a  
6 high level, particularly, you said working with --  
7 like what did you actually collect and identify when  
8 you got these data systems? So that way I'm not  
9 asking you questions for things that you're, like,  
10 "I didn't do"?

11 A Yeah. Yeah.

12 So, well, the way the program operated was  
13 we were surveying or learning through working with  
14 the various teams that we knew to have, you know,  
15 managed storage systems and we basically asked, you  
16 know, some simple questions about whether or not  
17 they could store user data or if there's a  
18 possibility that user data was stored in their  
19 systems.

20 And then from that point if they said yes,  
21 then we worked with them to understand how they  
22 classified user data and not user data. So it was a  
23 simple kind of binary check. It wasn't, you know,  
24 details of the data, but just user data, yes/no, for  
25 any given unit of storage; right?

1           So let's say it's a database, a relational  
2     database, does that table have user data? Can, you  
3     know, rows of this have user data? If they say yes,  
4     then how do you -- we worked with them to document  
5     how they classify that data and then configure it  
6     for deletion to meet the requirements and, in  
7     particular, the requirements of our user data  
8     deletion policy.

9           Q     When you asked --

10          A     And --

11                Yeah. Go ahead.

12          Q     I guess when you asked them questions, so  
13     like -- because if you look here, right, in  
14     Exhibit A, and we'll go through it in a lot more  
15     detail, there's 149 of these systems.

16                So you would email somebody and they would  
17     send you -- did you -- what kind of questions did  
18     you ask them? Like what -- because I don't want to  
19     ask you to do work you haven't done or for  
20     information you don't have. I think if I understand  
21     what you have, it may streamline the questions  
22     process for the systems.

23          A     Yeah. It would be basically, you know, a  
24     simple description of their system. Whether or not  
25     there's user data stored; if they maintain backups;

1 if they -- you know, what are the deletion options  
2 available if it's data. Things of this nature.

3 You know, I don't have -- it's been some  
4 time since I've directly worked with that. I  
5 started the program and the team, you know, took it  
6 forward and I haven't been reviewing the current  
7 intake form and kinda the questions that are asked  
8 there, but they're of that nature. Fairly  
9 high-level questions.

10 Q So then what information or artifacts? So,  
11 you know, what you -- like does it -- so for each of  
12 these systems, you -- I'm just trying to understand,  
13 like, would they tell you what user data they were  
14 getting or, like, what in- -- like, for example, so  
15 for all of -- I'll give you an example.

16 For the systems, and we'll start with one  
17 of them just as an example, right, and I'll scroll  
18 down here to -- and I received a spreadsheet with a  
19 letter that broke down what they were.

20 A Uh-huh.

21 Q And so if you look at ABP, it says AWS, it  
22 says files with salted and hashed Facebook users'  
23 PII are stored in this data store. The data store  
24 is used by the [REDACTED]

[REDACTED]



1 [REDACTED].

2 My first basic question, you may not know  
3 the answer to this, would you know if this user data  
4 is already contained in the social graph?

5 To me it seems like it's already there.  
6 It's redundant.

7 Is that -- would you know the answer to  
8 that question?

9 A No. I wouldn't know.

10 Q Okay. Would you be able to tell me like  
11 why was this system identified, the ABP, for having  
12 user data? Like, is it --

13 A We could -- we could trace back to see  
14 where we learned about the system. Given that it  
15 is --

16 Q Well, before we --

17 A Yeah. Go ahead.

18 Q No. Go ahead. Sorry.

19 A No. I was just going to say, you know, a  
20 clue there, being that it has AWS in parens  
21 indicates that it was through our AWS, like a  
22 specific AWS effort to identify all AWS systems.

23 So we knew of use of AWS, went to the team  
24 that manages that, and worked with them to get the  
25 list of those data systems.

1           Q     I mean, did you find out if any of the data  
2     in these systems, the user data you identify is  
3     actually -- is it shared with -- because, I mean,  
4     for some of these, you know, is the data being  
5     shared with third parties or -- and, I mean, whether  
6     it's readable or not readable, I don't care.

7                     What I want to know is, is the data in the  
8     system shared with third parties, like ABP?

9           A     Yeah. I don't know.

10          Q     For all of them or just that one?

11          A     So for some I will know because I have more  
12     familiarity with them.

13                     This one I'm not familiar with directly;  
14     but some that I have worked with, I would be able  
15     to -- I have a little more depth of knowledge and I  
16     could.

17          Q     All right. And then my other question is,  
18     is the data -- and these are the two real questions.

19                     Is the data that's collected and the user  
20     data, being very clear, the user data collected made  
21     accessible to other parts of Facebook that then put  
22     it into a product?

23                     And I'm not saying for a particular -- I'm  
24     just going to give you an illustrative example, and  
25     I'm not saying this actually happened for the

1 record. It's just to illustrate an idea; okay? I  
2 have no personal knowledge.

3 A All right.

4 Q But this data falls -- goes into your  
5 system, then it outputs to another system at  
6 Facebook that then uses some of that output and  
7 provides it to an off-platform ad network or  
8 something; right? Is the data being made accessible  
9 to another platform offering or solution that  
10 Facebook shared off, you know, with a third party,  
11 whether it's an ad network, whether it's with  
12 whoever? I mean, maybe ad network is a very touchy  
13 subject. Probably shouldn't use that one.

14 But the idea is that it's shared off  
15 platform, but the data goes into the system and then  
16 that data is processed or something's done with it  
17 and then shared with another entity maybe likely on  
18 this list, I would assume, that then outputs it.

19 A Yeah. You know, actually, I don't know.  
20 It's kind of like outside the scope of, like, what  
21 our strategy was for -- and ultimately the program  
22 was around ensuring that we could meet the  
23 requirements of the XT order.

24 Q Okay.

25 A Section 3, deletion of information when a

1 user initiates deletion.

2 We didn't get into how data is processed or  
3 these other kind of, you know, life cycle of the  
4 data.

5 Q Okay. Fine.

6 So I guess for the systems we do know,  
7 right, so we'll go through the list fairly quickly.

8 I mean, for AdFox, you know, the  
9 description is the delivery flow, which I have no --

10 Do you have any idea what that means?

11 A This is used in the delivery flow for  
12 targeting, ranking features.

13 Q Is that to customers? Is that an internal  
14 tool that you're using to manage in the ad network?

15 A Yeah. I don't know, actually. I'm not  
16 familiar with this system or kind of like the  
17 relationship that they're describing between  
18 Unicorn and so forth. Yeah, I haven't worked in  
19 this area to know what that means more than what's  
20 written.

21 Q Just so I understand, when it says "the  
22 social system," is that the AdFox system?

23 A I don't know. Yeah. Sorry.

24 Q Okay. So do you know who --

25 But you have a point of contact for each of

1 these systems that could theoretically -- if I  
2 think -- and that's a big if, in caps -- if I think  
3 it's necessary or appropriate to get the additional  
4 information if this is used or shared.

5 And just so I'm clear here, their user data  
6 would go into this system that -- and I'm not sure.  
7 Would this system then share the data off platform?

8 A I don't know.

9 Q You don't know?

10 A Yeah. I'm not sure how the data is  
11 processed.

12 Just, again, kind of the scope of where we  
13 limited our view of it was, you know, identify the  
14 system, right, so now we know that AdFox exists;  
15 identifying a point of contact to work with; and  
16 then asking them can the system store user data,  
17 which they indicated yes, that's why it's on this  
18 list; and then it was, okay, for user data stored in  
19 your system, how are you configuring it for  
20 deletion?

21 Q Very good.

22 A That's the limit of what we get into.

23 Q Yeah. It's a GDPR data map. So  
24 effectively you could manage your data privacy  
25 requirements or CCPA or whatever regulatory --

1 A Right.

2 Q -- framework you were looking in.

3 A Right.

4 Q In this list, we'll go through each of  
5 them, but if there are specific ones you know more  
6 about, please tell me. Because I read about  
7 AdLogger and let's be clear. All I'm interested in  
8 is really if the data is shared off of the platform  
9 or made accessible off of the platform. And when I  
10 read the description of AdLogger, it says it's a  
11 real-time service that performs a join, which I  
12 assume you mean a database join, between different  
13 event streams, which I interpret to be I'm not sure  
14 what, and features. And it says the core -- but the  
15 part of that that's kind of confusing, it says core  
16 of the training data preparation flow, which to me  
17 says is this an internal tool that you're just using  
18 to optimize for the ad network that you then  
19 aggregate out?

20 A Yeah. You know, I can -- I don't want to  
21 speculate. I don't actually know the details.

22 I think there are some hints there when  
23 they say like feature and so forth, but I don't want  
24 to infer because I -- I don't actually work with  
25 that system and don't have knowledge of it.

1 Q Okay. I appreciate it.

2 But you, again, have the spreadsheet  
3 response or the form response you got for that  
4 because --

5 A Right. We do have the point of contact, as  
6 you mentioned, for that system.

7 Q Yeah. All right.

8 Then for AdMarket database, I think this is  
9 a MySQL database, it's an internal system, but is  
10 that correct? I mean, what is it doing because --

11 A I mean, again, you know, unfortunately, I  
12 guess, my knowledge is --

13 Q No. It's okay. I much -- as crazy as this  
14 may sound, I much appreciate your repeating yourself  
15 and you can just say, "Same answer as last time," if  
16 you want to streamline it and it won't offend or  
17 upset me at all.

18 A Yeah.

19 Q All I'm trying to do is figure out. You  
20 submitted the declaration. So I wanted to check  
21 with you.

22 You don't have any further information  
23 available, right, that you could reference to check  
24 and give answers?

25 A Yeah. Not at hand. Certainly, you know,

1 this knowledge exists within the company, but it's  
2 not essentially, you know, a part of the work that I  
3 did or that I have access to.

4 Q I got it.

5 A Yeah.

6 Q Because it says -- and the reason why I  
7 ask, it says but also other critical consumers of  
8 the ads system Ads Delivery.

9 So, again, is this user data being shared  
10 or aggregated out of the Facebook platform into  
11 third parties or is this just Facebook's internal  
12 systems that it's using?

13 A Yeah.

14 Q Well, maybe --

15 A I don't --

16 Q Well, have you had a chance to look at all  
17 149 systems?

18 A Yeah. I mean, I do -- like, if you want me  
19 to call out a particular one, CDN is the one that  
20 I probably have the most knowledge of and  
21 understanding.

22 Q And CDN is?

23 A Yeah. It's the Content Delivery Network.

24 Q I'm just trying to scroll down.

25 A I think you passed it. I think I saw it go



1 by.

2 Q Sorry. I was using my finger. One sec.

3 Where would it be?

4 A It's a little bit higher.

5 Q CDN.

6 A CDN.

7 Q There it is.

8 A There it is above.

9 Q There it is. Okay.

10 A Kinda on the fold of the page. Yeah. It's  
11 on the list, yeah.

12 Q So the CDN says CDN is a team within  
13 Traffic dedicated to the performance and delivery of  
14 media content of all Facebook products.

15 How does user data play into that?

16 A So the way this CDN works is [REDACTED]

[REDACTED]  
[REDACTED]  
[REDACTED]  
[REDACTED]  
[REDACTED]  
[REDACTED]

22 Q Yeah.

23 A [REDACTED]  
[REDACTED]

25 Q Yeah. It gets delivered faster to the

1 user.

2 A Yeah. Exactly.

3 So, you know, [REDACTED]

■ [REDACTED] [REDACTED]

■ [REDACTED] [REDACTED]

■ [REDACTED]

■ [REDACTED] [REDACTED]

■ [REDACTED]

■ [REDACTED]

10 Q Right. And then --

11 A And so when --

12 Q Go ahead. Keep going.

13 A So then [REDACTED]

■ [REDACTED]

■ [REDACTED]

■ [REDACTED]

17 Q And so then [REDACTED]

■ [REDACTED]

■ [REDACTED]

■ [REDACTED]

■ [REDACTED] My question is, do you know

22 what my -- what's it called? Download. I forget

23 what the tool is called.

24 A Download your information.

25 Q DYI. Sorry. Yeah.

1           Is all of the information in a DYI the same  
2     information that's available in the CDN because  
3     those are my pictures and the DYI, as I understand  
4     it, has those photos? Is that correct?

5           A     I don't know. I don't know. I don't  
6     actually know how DYI operates or the systems that  
7     it covers.

8           But for the CDN, again, our scope was is  
9     there user data stored there. So the answer was  
10    yes. It would be like a picture of a cat.

11          Q     Yeah. Yeah.

12          A     That --

13          Q     It's --

14          A     -- is deleted.

15          Q     It's an operational storage of data that  
16     then is used to support the user experience. It's  
17     not --

18          A     Yeah.

19          Q     -- user data, per se, that's being shared  
20     off platform to a third party. It's being made  
21     accessible to my friend in Tel Aviv so that they can  
22     access the picture.

23          A     Right.

24          Q     Does any other third parties have access to  
25     the CDN data or is it just end users?

1 A Well --

2 Q Like does Facebook make CDN data --

3 A Yeah. No. It's available publicly. So,  
4 you know, I guess there's not a special channel that  
5 I'm aware of that would be for that. But this is  
6 really, again, kind of outside of my area of  
7 expertise. The knowledge I have of CDN is the  
8 public-facing --

9 Q Okay.

10 A -- use --

11 Q But it's used --

12 A -- like in Tel Aviv as you described.

13 Q And it's used to support the function of  
14 sharing graphics, videos, or photos through the  
15 Facebook platform --

16 A Uh-huh.

17 Q -- that a user uploads or content they  
18 upload and it's made publicly. However the privacy  
19 settings are set out wherever, they get wrapped  
20 around the image in the CDN at some point?

21 A Right.

22 Q All right. So then how about Calypso?  
23 Nothing? Same as before?

24 A Yeah. Same as before.

25 Q When they said user data comes into it,

1 nobody -- the problem is is the plan is --

2 So this looks like a new system. Is that  
3 an accurate --

4 A I don't know the age of the system.

5 Q Because it says small set of initial  
6 customers. The plan is to build out the  
7 Presto-based solution as soon as possible and  
8 transform it onto the platform.

9 All right. And then Callisto, you know,  
10 again, was created -- the part I'm interested in, it  
11 says it was created -- and I guess this was provided  
12 to you, so you're sure. But it says it was created  
13 in order to migrate from Prometheus, which I'm  
14 pretty sure I know what they're talking about, but  
15 let's just assume on the MySQL with RocksDB storage  
16 platform that Facebook uses.

17 A Yeah. I know a little bit --

18 Q That's an internal.

19 A -- about that one. Yeah.

20 Q So that's an internal system; right?

21 A It backs the messaging service. So  
22 actually, the messages you send, you know, through  
23 Facebook are -- this is like the application layer  
24 to that service.

25 Q So if it's just --

1 A That application.

2 Q Yeah. So only Facebook internal services  
3 have access to the application layer or do  
4 third-party applications have access to it as well?

5 A No. That one is actually --

6 Q Which one is right?

7 A So only internal applications and products  
8 have access to this system.

9 Q All right.

10 A And what I just want to be clear about  
11 that.

12 So the messaging service is exposed  
13 externally, right, like the actual messaging  
14 takeout, you know, delivery.

15 Q Yeah.

16 A That does not -- that has access to this.  
17 But beyond that, you know, users of those messaging  
18 applications can't go directly to Callisto.

19 Q Got it.

20 And then for Chronos, what user data is  
21 going in through Chronos?

22 A Yeah. I have a little bit of knowledge  
23 there.

24 So Chronos is, you know, a scheduler. The  
25 reason that it has been identified as can store user

1 data is that you could schedule a job with user data  
2 in the description.

3 Q But that's, again, for operations internal  
4 to Facebook --

5 A It's completely internal.

6 Q -- or is it --

7 A No. It's completely internal.

8 Q All right. So then Chronos --

9 So you're talking about you could schedule  
10 an internal job for whatever on a database cleanup.  
11 The Chronos will run on the database and then it  
12 will report -- you'll schedule a database cleanup.  
13 Whatever the activity is, that runs only internal to  
14 Facebook; right?

15 A That's correct.

16 Q Yeah.

17 A I do have knowledge of Chronos to that  
18 degree. That's the -- pretty much the depth that I  
19 have.

20 Q That's all I need to know, honestly. So  
21 that's perfect.

22 Cocoon.

23 A Yeah. I'm not aware of that.

24 Q Again, inventory estimation. I assume  
25 that's an internal tool, but I guess you're not the

1 right person to ask, but I didn't -- because it says  
2 it's mainly used by advertisers trying to target  
3 their ads to the right audience. So I didn't get  
4 how user data fed into that and if it's being  
5 made access- --

6 Is Cocoon being made accessible to third  
7 parties, the data in there? It's user data really  
8 what I'm interested in being made accessible, but I  
9 guess you may not be the right person for that one.

10 A Yeah.

11 Q For Coefficient. I heard a lot about  
12 Coefficient actually very recently. But  
13 Coefficient, does that have user data?

14 A By the fact that it's on this list, they  
15 did indicate that it has -- it can store user data.

16 And just to make a qualification there,  
17 this is where the questionnaire was is it possible  
18 for a product or an application using the system to  
19 store user data there; not necessarily that they  
20 are, but that they could.

21 Q That would have been a good thing to put in  
22 your declaration, that they can store the data, but  
23 they don't necessarily.

24 So you don't know for sure that they do  
25 store user data there, but it can store user data?



1           A     Right. It's possible that user data is in  
2     that system, and so we take the position that we  
3     must secure it, you know.

4           Q     I got it. I get the policy driving piece  
5     of it.

6                     And then for Co- -- and then just to go  
7     back up really quickly just to run through the As.  
8     I had a couple questions.

9                     Again, AdMarketDB, it says it's the  
10    backbone for the Core Ads system.

11                    That is not exposed to third parties or  
12    made accessible to third parties or is it?

13           A     I don't know. I don't know much about the  
14    ads system.

15           Q     All right. And then so I won't ask  
16    about it.

17                    I assume same answer for Ads Raw Storage?

18           A     Yeah.

19           Q     Okay.

20           A     No, I don't recall anything more on that.

21           Q     All right. For Akkio locality.

22           A     I do have a little bit more knowledge of  
23    that one.

24           Q     What can you tell me about this one?

25           A     So Akkio is -- the purpose of this is to --

1 if you have data that really only you are accessing,  
2 right, through our products, this --

3 Q When you say "you," do you mean a user or  
4 do you mean a Facebook employee?

5 A A user.

6 Q I don't want to --

7 A An outside -- yeah. An outside non-Meta  
8 employee who is using our product; right?

9 So, you know, my mother using the Facebook  
10 app, there's certain information or certain data  
11 that might only be for, like, her configuration  
12 settings; right? It's not like publicly available.  
13 She is the only person that would look at that data.  
14 And so what would happen is it determines,  
15 it detects where in the world she's accessing  
16 the prod- -- you know, the product is accessing the  
17 data from and it will relocate that data to a data  
18 center closer to her. So if she's on the East Coast  
19 and moves to the West Coast, after some time it  
20 figures out, oh, we should move her data.

21 Q I got it.

22 And, again, it's moving her data, which is  
23 what --

24 When you say "the user's data," that's the  
25 data available through the social graph or the data

1 is broader than that? Like is it all of their data?

2 A No. It's product specific, so -- or  
3 feature -- it's actually probably more feature  
4 specific. So case by case, they would configure the  
5 system to provide that capability.

6 So that's about -- that's the depth that I  
7 have knowledge of is what it's -- the feature is  
8 providing, but I don't really have the use cases or,  
9 you know, other information.

10 Q Did they provide use cases when they  
11 submitted it to you? Did they give you use cases in  
12 their submission?

13 A Yeah. I'm sure it was discussed. I don't  
14 think we -- you know, this is -- it isn't a part of  
15 the questionnaire or what have you. But when we  
16 actually meet to understand what they're doing, most  
17 likely, they would have described, you know, a  
18 scenario.

19 Q You have some description of what -- for  
20 each of these systems, just so I get the process,  
21 you guys would send the questionnaire, they'd send  
22 you the response, you read the response, probably  
23 talk to legal, talk to whomever, make a decision  
24 and then schedule a follow-up meeting, if deemed  
25 appropriate, to understand the use cases for the

1 systems and that your team did -- not you, per se,  
2 but when I say "you," I mean the royal you -- the  
3 team that you were working with would go and  
4 interview and get use cases or whatever for that  
5 particular system and that's how that was handled.  
6 Is that accurate?

7 A So it would be that it's possible that use  
8 cases would come up and I think in this particular  
9 case I was involved with it. It's been some time.  
10 So my memory isn't fresh about the details. But I  
11 can imagine that we did just talk about use cases.

12 It's not a part of the process, though.  
13 It's not like we go and say explicitly as part of  
14 the engagement, "Okay, let's discuss use cases." It  
15 would be an extension of discussing, "Okay, how are  
16 you capturing the classification of this data and  
17 how are you configuring it for deletion?" And that  
18 could lead to, but is not required to lead to a  
19 discussion on use cases.

20 Q Okay. And so for Ads Online Storage,  
21 that's only available or made accessible to internal  
22 Facebook or external as well?

23 A I don't know on that one.

24 Q You don't know on that one. Okay.

25 Okay. Archival.

1 A I have a little knowledge --

2 Q It's a backup tool.

3 A -- on this.

4 Yeah. This would be internal only.

5 It's --

6 Q Got it. Okay. Great.

7 A -- for providing storage.

8 Q That's all I need. Internal only. We're  
9 good.

10 A Okay.

11 Q And it's the backup. So it's not -- yeah.  
12 Artillery, is that how I should say it?

13 A Artillery, yeah.

14 Q Is that a development tool or --

15 A It's -- yeah. I actually don't have more  
16 knowledge than what's kinda written in that  
17 description.

18 Q Okay. But, again, is it only available to  
19 Facebook internally, the tool set?

20 A I don't --

21 Q Because it can be read maybe as a tool set  
22 available to app developers or is it a tool set only  
23 internal to Facebook?

24 A I don't know. Actually, I'm not sure.

25 Q But in your system of the data you

1 collected, it would have more information about the  
2 system at some level than these three lines or five  
3 lines here?

4 A It will have more information, but again,  
5 oriented around the classification of the data and  
6 the configuration for deletion when initiated by a  
7 user, not necessarily how the data is processed or,  
8 you know, what the actual data is.

9 Q Right.

10 A And we just took kind of like that binary  
11 yes/no user data. Not anything beyond that.

12 Q And Assistant User Memory, is that -- first  
13 question, is that available or accessible to  
14 external parties, third parties, that data?

15 A I don't -- I don't know actually on this  
16 one.

17 Q All right. And then, I assume, Async is  
18 like Chronos, but I didn't want to -- can you  
19 confirm that or no?

20 A Yeah. So I do have a little bit more  
21 knowledge of Async.

22 Async is a system that runs [REDACTED]  
[REDACTED]  
[REDACTED]  
[REDACTED]

25 Q I get it.

1 But it's not available to anybody off  
2 platform?

3 A Correct. It's just -- yeah. It's just a  
4 way -- it executes code that is running on our  
5 platform.

6 Q I got it.

7 It's got user data that's used for your  
8 web-based applications for Meta and you use it  
9 internally. So you identified it as a user system  
10 because it has user data for purposes of the other  
11 issue you were addressing; but at a functional  
12 level, that data stays on Facebook, isn't shared  
13 with anybody. It's used internally at Facebook for  
14 those jobs?

15 A Yeah. Just to clarify, I think what you're  
16 saying is correct, but just to clarify.

17 It is like -- [REDACTED], it is the  
18 same as that. It's just that for like, let's say,  
19 you make a request external to [REDACTED], user of the  
20 Facebook product, and if you make a selection of  
21 some option that doesn't have to [REDACTED]  
[REDACTED] and the [REDACTED]  
[REDACTED] would be  
24 able to take that and then make -- you know, kind of  
25 [REDACTED]

1 [REDACTED] like where --

2 Q Right.

3 But that's used by Facebook, that's used by  
4 you guys internally [REDACTED]; right?

5 A Yeah. Yeah. Yeah. But I just want to  
6 make sure that's not --

7 Q We can go on.

8 A Yeah.

9 Q But it's used to [REDACTED] right?

10 A Yeah.

11 Q It's used for -- it's for --

12 Okay. And then, again, Bars. I don't mean  
13 to sound ignorant. I didn't know what the Bars app  
14 even was.

15 A Yeah. I -- so I can't -- I don't know the  
16 Bars app.

17 But what I can tell you from this is that  
18 NPE is new product experience. So this would have  
19 been a new product in kind of prototype mode  
20 probably. But that's where, you know, a zero to one  
21 new product experience app, expand the acronym.

22 Q Got it.

23 And new product experience in the last  
24 year, two years?

25 A I don't know any kind of the -- so



1 basically what we would have with that is if it is a  
2 viable application, then it would move to more  
3 formal production. Like it would actually become  
4 internally, you know, outside of that system or  
5 the -- it would move out of the new product  
6 experience to a more formal development team for  
7 long-term development.

8 Does that make sense?

9 Q Yeah. Of course.

10 So then Beat Saber, Multiplay and  
11 Analytics, are those internal tools? Because when  
12 I read it, it's not clear to me. Is that made  
13 accessible to third parties or is this internal tool  
14 sets that then you use to aggregate through the  
15 Facebook platform?

16 A I believe this was an acquisition of a  
17 gaming system or gaming --

18 Q Okay.

19 A -- for virtual, virtual reality.

20 Q Oh, okay. So you're not sure, one way or  
21 the other, if that made -- the data there is made  
22 accessible off Facebook's platform?

23 A Right. I don't know.

24 Q Like if I have my high score, users can get  
25 it or it's made available.

1           A     Okay. What I explained, I mean, that's the  
2     depth of understanding I have.

3           Q     Okay. So then BeeST and local name space.

4           A     Those are internal.

5           Q     That's completely internal?

6           A     Yeah.

7           Q     And then Bento Next.

8           A     It's also an internal tool.

9           Q     And the results of that tool, it says "and  
10    share them with others." That's to share internal  
11    to Facebook, right --

12          A     To other --

13          Q     -- or is that to share with others?

14          A     To other employees of Facebook. Basically  
15    you can write these kind of data processing scripts  
16    to generate reports and so forth.

17          Q     Got it.

18                And BigBox VR. It says "M & A." I'm  
19    assuming --

20          A     That's an indication --

21          Q     Well --

22          A     The M & A tag would indicate that it was  
23    through our mergers and acquisitions. So this was  
24    an acquired company that we then, you know, do the  
25    same process; right?

1 Q But it's not part -- it is part of the VR  
2 AR system or not the Facebook web platform?

3 A Yeah. This would be -- the VR indicates to  
4 me that this was a part of the virtual reality  
5 services.

6 Q Yeah. Right.

7 So that it's not the Facebook website or  
8 social experience or mobile app as we understand it?

9 A Right.

10 Q All right. We covered Callisto and  
11 Calypso. We covered CDN and Chronos. I asked you  
12 about Cocoon. Coefficient.

13 Collab. It says powers the Collab  
14 application. It says Collab is a standalone app for  
15 creating music video content collaboratively. The  
16 data stored on our system is mostly user data.

17 Do you have any idea if that user data is  
18 accessible to third parties off platform?

19 A I do not know.

20 Q Configerator. I don't know if that's a  
21 play on words or Configerator?

22 A It's got two pronunciations, which one  
23 you've got correct. Configerator and the other is  
24 called Configerator. It's an internal configuration  
25 for applications and so forth. So it's internal.

1 Q Got it.

2 And then Connectivity Testing Framework,  
3 that's hardware. So I assume that's not related to  
4 the social piece, but is that correct, that  
5 assumption?

6 A I don't know. Yeah. Yeah. I would be  
7 speculating, but it does seem a reasonable  
8 inference.

9 Q All right. And then Contributor Store, aka  
10 Zippy, is used to store contributors. I didn't know  
11 how you define what's a contributor.

12 A Yeah.

13 Q Do you know, by chance?

14 A I don't know.

15 Q And then my question is --

16 A And just for --

17 Q My real --

18 A Yeah. Sorry. Go ahead.

19 Q No. I was just going to ask, does this  
20 data go off the system anyways or is this all  
21 internal?

22 A Yeah. I don't know that.

23 Just to clarify the use of parens here. So  
24 this is indicating that [REDACTED]

[REDACTED]

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

Q Okay. So didn't know that's how it's connected. I was wondering because that's in a couple places, which makes a lot more sense then. It's like, Why are you repeating systems? But now I get it.

Okay. Cortext, a digital asset management tool or referred to as DAM.

Is that an internal tool or an external tool here? Like can a user -- is this data, is it accessible off Facebook's platform or shared off the platform?

A Yeah. I don't know.

MR. GARRIE: Did we lose the court reporter?

THE COURT REPORTER: No. I'm here.

BY MR. GARRIE:

Q So returning back.

Cortext, again, I --

A A digital asset management system.

Q Yeah. Okay. All right.

Cosco. I assume this is an internal tool accessible only in user data, but it's only user

1 data for the operations of Facebook.

2 Is that an accurate assumption?

3 A Yeah. My knowledge of the system is that  
4 it is for internal reporting and, again, it's kinda  
5 like [REDACTED] in a way for internal tools. [REDACTED]

6 [REDACTED]

7 [REDACTED]

8 [REDACTED].

9 Q That's what I thought.

10 And Cubrick is the same sort of genre or is  
11 it different in that it's memory in real time, but  
12 same idea, internal?

13 A That's -- to the best of my knowledge,  
14 that's correct.

15 Q This is a big terabyte. 10-terabyte tables  
16 are, like, some big old tables.

17 That's what it means; right? Up to 10-plus  
18 terabytes; right? That's for the tables, queries on  
19 those tables; right?

20 A Yeah.

21 MR. GARRIE: Oh, and this transcript needs  
22 to be marked confidential for the court reporter.

23 THE COURT REPORTER: Okay.

24 MR. GARRIE: And then we'll work it out  
25 from there with the parties afterwards just for

1 purposes of keeping it all flowing.

2 Q CWS, Core Workflow Services. Again, can  
3 you confirm, is this internal data or is this made  
4 accessible to third parties?

5 A Yeah. If I remember correctly, this is  
6 actually a backing component of Async.

7 Q All right.

8 A So, again, kind of like that internal  
9 workflow management.

10 Q Yeah.

11 The reason why I'm asking is, right,  
12 because it says it's designed to let people and I  
13 just didn't know if those people were the third  
14 parties or --

15 A No.

16 Q -- they're Facebook employees.

17 A Those would be Facebook employees.

18 Q And then the Daiquery. I'm not sure how to  
19 pronounce that one.

20 A Yeah.

21 Q So I'm not even going to try.

22 A Yeah. That's just -- it's a UI tool for  
23 running queries on back-end databases used  
24 internally.

25 Q Internally.

1 All right. It's actually much more  
2 productive than I thought it was going to be. So  
3 we're doing well.

4 DAPL, D-A-P-L, it says it's a back-end  
5 platform to build audience segments.

6 My question is, is that accessible to third  
7 parties or when you say "back-end," is it only back  
8 end to Facebook?

9 A I'm not familiar with that system.

10 Q And we'll just say Dark. It says [REDACTED]

[REDACTED]

12 I assume that's an internal tool, but can  
13 you confirm that?

14 A Yes. That's [REDACTED]

[REDACTED]

16 Q And so Gluster FS, Dataswarm is built on  
17 Gluster FS.

18 A That's correct. It's, yeah, the Dataswarm  
19 instance on Gluster FS. Yeah.

20 Q And is this an internal tool set that's  
21 used? Because I wasn't --

22 A That's correct.

23 Q -- exactly sure who has access to it.

24 A Yeah. This is an internal-only tool.

25 Yeah.



1 Q And then Deld? Deld? I don't know to --

2 A Deld is how they refer to it.

3 Q Deld. Thank you.

4 Is Deld a --

5 A I do have some -- yeah. I can tell you a  
6 little bit about this one. I do have some knowledge  
7 about it.

8 It's basically the deletion technology for  
9 WhatsApp. So when the user deletes data on  
10 WhatsApp, this is the system that then deletes that  
11 from the underlying storage.

12 Q Yeah. But the user has no direct access to  
13 this; right? It's not --

14 A No. No. They don't have direct access to  
15 this. The access would be through -- it's the same  
16 as if Facebook -- you know, the -- the Facebook  
17 product, someone deletes, and then it goes through  
18 deletion.

19 Q All right.

20 A This is specific to WhatsApp.

21 Q Yeah. Okay.

22 A And you don't have -- you don't have access  
23 to it. It runs on the back end.

24 Q And then Deletion Framework.

25 A Yeah. So that's our --

1 Q What's this www?

2 A So that's for [REDACTED] [REDACTED]

3 [REDACTED] And this is the  
4 central [REDACTED] that may have been  
5 discussed in other conversations that will delete  
6 data when someone requests it to be deleted through  
7 our application.

8 Q Got it.

9 But, again, it's only available -- this is  
10 data -- the data in the platform --

11 A It's only internal.

12 Q -- is only --

13 A It's only internal.

14 Q Got it.

15 And Delivery Insights. I didn't know what  
16 signals meant. So...

17 A Yeah. I'm not -- I'm not familiar with how  
18 this data is used.

19 Q Because I wasn't sure if that means that  
20 they're taking --

21 Do you know what kind of user data? Is  
22 that individual user data going into this system or  
23 anonymized user data? Did they tell you what kind  
24 of user data they're collecting for these systems,  
25 by chance?

1           A     I don't know. No. I didn't work with  
2     Delivery Insights.

3           Q     And is this Facebook Delivery Insights?  
4     Facebook, not the virtual. Not, you know, virtual.  
5     It's specific to -- do you have any idea?

6           A     No. I don't have any idea.

7           Q     Dependency Manager Service. I thought it  
8     was only internal, but I just want -- it's only  
9     internal access. The data never goes off platform  
10    here, user data?

11          A     Yeah. I can't really answer that. I'd be  
12    speculating on that one. I'm not familiar with it.

13          Q     DevDB is --

14          A     Yeah. I am familiar with Dev and it is  
15    internal only.

16          Q     And the data in there -- okay.

17                And it says DevOps Investigation Tool  
18    (DIT).

19                Is it built on DIT or is it called  
20    DevOps Investigation Tool in this case?

21          A     Yeah. I think this is just its, you know,  
22    acronym.

23          Q     Got it.

24                This is available only internally? Because  
25    what does SEV stand for?

1           A       Site event and we refer -- we internally  
2       call this SEV. So this would be when, you know,  
3       something breaks in production. So there's a site  
4       event, and then from this description this is what  
5       the DevOps teams use to investigate.

6           Q       But it's only available and accessible to  
7       the DevOps team?

8           A       Right. I -- I don't know if I see there  
9       that explicitly stated. So I don't -- I can't  
10      confirm, but it is an internal tool. Site events  
11      are internal, you know.

12          Q       We can leave it there.

13          A       Yeah.

14          Q       What's a DiGraph?

15          A       So just going by the description,  
16      Data Infrastructure Graph is a compute engine for  
17      large-scale graph analytics.

18                 Yeah. You know, I'd be speculating and  
19      just kind of inferring from what's written there. I  
20      don't have experience or knowledge of that system.

21          Q       And, more importantly, is it available or  
22      is it open to third parties or not?

23          A       Yeah. I'm not sure.

24          Q       And then Downpour Interactive, that's  
25      through M & A?

1 A Yeah.

2 Q Is this a game on Facebook's platform or is  
3 it separate from Facebook's platform?

4 A I don't know, actually, but I imagine it's  
5 actually our reality lab. That's where most of  
6 these M & As are. It's for games that are acquired  
7 for --

8 Q All right.

9 A -- this.

10 Q Okay. Got it. That's very helpful. I  
11 kinda Googled it. I just couldn't figure out if it  
12 was integrated into the larger platform.

13 Okay. So then Dumbo --

14 A Uh-huh.

15 Q -- like the elephant. Is that internal?

16 A It's internal storage. Like a backup  
17 system.

18 Q And no third parties have access to that  
19 and the data stays there; right?

20 A That's correct.

21 Q Okay. Then we have dx/Models Framework.

22 A Yeah. Yeah. So I don't have knowledge of  
23 this, but I can tell -- like I can tell you what  
24 Novi is referring to. That's the cryptocurrency  
25 solution we're developing. So this would be part of

1     that system.

2           Q     Okay. But that's done in the crypto space  
3     and not in the Facebook platform itself?

4           A     That's correct.

5           Q     So when it says -- because the reason I'm  
6     even -- it says "risk and compliance assessments."

7                     That's available only internal to Facebook.  
8     Is that an accurate assumption?

9           A     I can't confirm that.

10          Q     I wasn't exactly sure how user data fit  
11     into that. So I was hoping you could clear it up.

12                     All right. EE Storage.

13          A     Well, just to offer. You know, Novi being  
14     a cryptocurrency wallet service. That would be  
15     where customers or, you know, external people using  
16     the Novi cryptocurrency services would be user data.

17          Q     So then this is a system that sits on top  
18     of the cryptocurrency wallet. So it's used for  
19     compliance and risk assessments as an internal tool  
20     for regulatory reporting maybe. Okay. But it's  
21     solely related to the crypto platform?

22          A     That's correct. That's the -- the Novi is  
23     the indicator for that, yeah.

24          Q     Okay. And EE Storage?

25          A     Uh-huh. This is an internal only.

1 Basically it operates like shared -- on a shared  
2 network drive.

3 Q Okay. And ElasticSearch?

4 A Yeah. So I can't really speak to the use  
5 of that and accessibility.

6 Q Sorry. I thought I had -- one of my kid's  
7 teachers might have had Omicron. So I was just  
8 getting an update.

9 All right. EverStore. Oh, sorry.

10 ElasticSearch is internal. I had a  
11 question about ElasticSearch and I -- ElasticSearch  
12 and some of these internal tools, they get data,  
13 right, and they process it. The output from that  
14 tool doesn't ever leave. Is it output to a  
15 different Facebook system?

16 A I don't know.

17 Q You don't know. Okay.

18 EverStore. What's DR compliant. Is that  
19 disaster recovery compliant?

20 A That's correct. That's correct. [REDACTED]

[REDACTED] So this is where the photos and  
22 videos primarily are stored.

23 Q And that's an internal tool. So that's  
24 user data, Facebook. This is where Facebook  
25 internally stores user data?

1           A     Right. That's where we -- internal and  
2     accessible only by the applications that run within  
3     Facebook infrastructure.

4           Q     F3. Facebook Feature Framework.

5           A     Yeah. This is referring to machine  
6     learning features.

7                     Beyond that, I don't have --

8           Q     How does user data relate to this?

9           A     So I -- the machine learning will operate  
10    on, you know, the history of user data.

11          Q     But it's only machine learning internal.  
12                     What I'm basically saying is Facebook, I'm  
13    a user of Facebook. I give you my data and it's  
14    only -- does the F, the feature framework, involve  
15    the user data being available to third parties or  
16    people off platform?

17          A     No. That would not be. Yeah.

18                     This case I have some familiarity with and  
19    it's used internally by the machine learning systems  
20    that then -- you know, an example of machine  
21    learning would be making recommendation on what  
22    groups you want to join or something like that based  
23    on --

24          Q     So then the output --

25          A     -- your behavior on the system. Yeah.



1 Q Well, where is it? Does that output then  
2 feed into another system?

3 A It would, yes. This is the -- the feature  
4 framework, this is providing a way of managing the  
5 machine learning features. So the machine learning  
6 system would use this. Internal only.

7 Q Internal only. Okay. I think I get that.  
8 All right. Fastpathdb.

9 A This one actually is externally accessible.  
10 I do have --

11 Q Okay.

12 A It's a message for when you're sending  
13 messages by messenger, and so it's temporary. You  
14 know, it's a receiving queue that then would  
15 interface with a back-end system like the ones we --  
16 Callisto, what have you.

17 Q So my question is --

18 A This --

19 Q Go ahead. Keep going.

20 A I was just going to say that this is the  
21 messaging interface exposed to the outside.

22 Q And then when you say "exposed to the  
23 outside," can you be more clear?

24 So if I -- let's say I write an app that  
25 uses messaging functionality. Is that even

1 possible? Is that where --

2 A No.

3 Q Because this says queue, persistent queue.

4 A Yeah. Yeah. I do -- again, I actually  
5 have a little bit more depth of knowledge on this  
6 one. I did some work with this as a messenger user,  
7 the application, the messenger application.

8 So it's directly coupled with use of  
9 Facebook products, messenger. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

15 A What's that?

16 Q So there's no data --

17 How old is the data in the system?

18 24 hours? Instantaneous? Like how old would user  
19 data actually be in the system?

20 A Yeah. I don't know, but it would -- you  
21 know, specifically, but it would be, I think, in  
22 line with what you're describing. Very short term.

23 Q So it's internal tool management,  
24 basically.

25 So when you say it's exposed to the

1 outside, just so I understand, that means it's  
2 exposed to Facebook messaging, like a consumer using  
3 Facebook messaging. I mean, it's not exposed to,  
4 like, a third-party app developer that builds  
5 something. They can't -- a third-party app  
6 developer doesn't have the ability to access this  
7 information or does it?

8 A For the -- what I worked on with this  
9 product or, you know, where I was involved with this  
10 and through the deletion program that we're mainly  
11 discussing now where the inventory came from, the  
12 messenger application, yeah, is the one running on  
13 your phone or your laptop would be the one that  
14 interfaces with this. It's not like a general --

15 Q Just that?

16 A -- queue for external developers to access.  
17 It's only for that messenger application.

18 Q Okay. So it's an internal tool Facebook  
19 messaging uses to facilitate messaging for an  
20 individual user using some data structure --

21 A Yeah.

22 Q -- to manage tools?

23 A I think that that's a fair description. I  
24 just want to make sure I'm not misrepresenting it in  
25 terms of, like, the internal tool part.

1 Q I said "tool." When I mean internal tool,  
2 I mean internal to Facebook, meaning that only  
3 Facebook messaging is the tool. Only Facebook  
4 messaging application can access this database,  
5 Fastpathdb.

6 If I build a Daniel Garrie data messaging  
7 tool, I can't access this database; right?

8 A That's correct.

9 Q Okay. And if I build an app, I can't  
10 access it as well. Only Facebook messaging can  
11 access this dataset; right?

12 A That is -- to the best of my knowledge,  
13 that is correct.

14 Q Okay. Fbcode, I assume, is just a [REDACTED]  
[REDACTED] and it has nothing to do with  
16 users.

17 I didn't get how user data, again, fits  
18 into this.

19 A Yeah. So this is a case where it is not --  
20 you know, [REDACTED]

[REDACTED]  
[REDACTED]  
[REDACTED]  
[REDACTED] -- what's the word?  
25 I'm losing the term.

1 Q [REDACTED]

2 A Yeah, but it's actually a different term  
3 that I was looking for.

4 But basically, you know, [REDACTED]  
[REDACTED],  
6 as you're describing, used by -- only for internal  
7 use.

8 Q So I guess my question is, no off-platform  
9 developers have access to this repository?

10 A Correct. And user data is -- should not be  
11 stored there.

12 Q Yeah. The bottom line is, it's only for  
13 you guys?

14 A Yeah. That's correct.

15 Q Okay. FBLearn.

16 A Uh-huh. I don't have a lot of knowledge,  
17 other than to say that this is a part of the machine  
18 learning artificial intelligence systems.

19 Q And the reason I'm even asking, it says all  
20 Facebook engineers. Is that limited just to that  
21 subset or can other engineers off platform gain  
22 access to the AI platform?

23 A I actually don't know the answer to that.

24 Q All right. And, again, built on Ziply.

25 A Yeah.

1 Q This isn't available to a third party --

2 A That's correct.

3 Q -- or is it?

4 A It's not, no. This is -- it is not. I do  
5 have familiarity with this system as well a little  
6 bit.

7 So this is users sending notifications to  
8 devices. So it provides those mappings of your  
9 account to device IDs and like, you know, being able  
10 to get to your phone to push a notification to your  
11 phone. So it's only used internally for that  
12 purpose of delivering notifications generated by our  
13 applications.

14 Q Okay. Got it.

15 And then FB Package.

16 A Yeah. That's just --

17 Q That's just binary?

18 A Yeah. Exactly.

19 Q How does user data -- well, I don't need to  
20 know.

21 The point is that this is only on platform.  
22 Never goes anywhere. Is that accurate?

23 A That's correct.

24 Q And it's never made accessible to anybody,  
25 but Facebook individuals?

1 A That's correct.

2 Q So FOQS.

3 A Yeah. This is another part of the Async  
4 system.

5 Q My only question is, it's only made  
6 accessible -- so it has user data, but is it made  
7 accessible to third parties, whether shared or off  
8 platform?

9 A Yeah. No. This is internal.

10 Q And then Forecast, that's a new product;  
11 right? NPE?

12 A Yeah. Yeah.

13 Q And I actually went and looked at that and  
14 I don't fully get the user data piece there.

15 A Yeah. I don't. I think actually you  
16 probably have more knowledge of it than I do, of  
17 that particular one.

18 Q All right. Now, ForgETS or ForgETS.

19 A Yeah.

20 Q I didn't understand this one.

21 My first question is, is it accessible off  
22 platform, the data in the system --

23 A No.

24 Q -- the user data in here?

25 A No. Not -- but, to clarify, I want to make

1 sure I'm not misrepresenting this.

2 Yeah. There is no external access to that  
3 system. So it would only be through applications  
4 running on Facebook infrastructure that could access  
5 that data.

6 Q Internal. But let's say a Facebook mobile  
7 app, that has access to this data infrastructure.  
8 But Daniel's Facebook mobile app, would I have  
9 access to that infrastructure?

10 A No. No. And this would be -- this is  
11 running behind the WhatsApp application.

12 Q Oh, okay. Then it makes a little more  
13 sense.

14 All right. Again, I didn't understand this  
15 FORT or FRL Research running on AWS and this is  
16 external research initiatives.

17 A Yeah. I don't know.

18 Q Is the user data made --  
19 You don't know?

20 A Yeah. I'm not familiar with that system.

21 Q And do you know what Jupyter --  
22 I don't even know how to say that.  
23 Jupyter.

24 A Jupyter.

25 Q Jupyter, with a Y. JupyterHub deployment,



1 what is that?

2 A My knowledge of Jupyter is, this is a  
3 public -- like, this is not a Facebook technology.  
4 It's for writing code in Python, I believe is the  
5 framework.

6 Q Yeah.

7 A Python notebooks.

8 Q Okay. Then Goals Tool, is that internal HR  
9 or is that like people that have life goals?

10 It says for teams to track goals  
11 information. I didn't know if that was Facebook app  
12 has goals or I didn't know.

13 Is this an internal tool?

14 A I don't know, actually. It sounds like it,  
15 but I don't know.

16 Q Okay. I'm going to agree with you on that  
17 one. Sounds like it.

18 The Hive. I've heard lots about the Hive,  
19 but maybe you can explain to me the Hive a little  
20 from your perspective.

21 A Yeah. That's our data warehouse. So, you  
22 know, we consider it the off-line storage facility.

23 Q Who has access to the Hive?

24 A Internal only.

25 Q Do third parties?

1           A       Yeah. This would not be directly  
2       accessible to third parties.

3           Q       Can a third-party app developer, though,  
4       access data stored in the Hive?

5           A       Not to my knowledge.

6           Q       Because it says an engine that executes  
7       queries written in FB-HiveQL and the Hive Query  
8       Language. So I didn't know who was actually --  
9       because it looked to me -- but you're not sure or  
10      you're sure?

11          A       These -- I cannot definitively confirm, but  
12      my knowledge of these, the data warehouse, it is,  
13      you know, for off-line storage and it -- the large  
14      number of use cases are related to [REDACTED]

[REDACTED]  
16          Q       Okay. And you don't know with any  
17      specificity what user data sits in any of these  
18      systems, right, just like you stated --

19          A       That's right.

20          Q       -- at the beginning?

21          A       That's correct.

22          Q       Hive Risk, I assume, is a -- it's only made  
23      accessible internally, the user data?

24          A       I can't confirm that. I'm not familiar  
25      with that system.

1 Q All right. Hotline/Fireside. What's it  
2 mean when there's a slash? Does that mean anything?

3 A Nothing in terms of how we've been tracking  
4 it. That was how it was given to us. So it must  
5 mean something within a team.

6 Q Got it.

7 Is this, you know -- is any of the data in  
8 this system discoverable off Facebook's platform or  
9 available or accessible?

10 A I do not know.

11 Q HulkDB?

12 A I am not familiar with this.

13 Q Didn't think you were, but it likely  
14 wouldn't be accessible? Facebook doesn't share any  
15 of the data in the HulkDB or make it accessible;  
16 right?

17 A I don't know. I actually am not familiar  
18 with this system.

19 Q You don't know?

20 A Yeah.

21 Q IRC. Is that the chat tool?

22 A That's correct. It's an internal instance  
23 of that. So it's used internal only.

24 Q Lacis. World AR. Is this augmented  
25 reality? I just wasn't sure.

1           A       That would be my understanding. I'm not  
2 familiar with the system, though, or World AR  
3 directly.

4           Q       All right. And then Laser.

5           A       Uh-huh. I don't know on Laser to the  
6 questions that you've been asking.

7           Q       Is it accessible or available to third  
8 parties or shared?

9           A       I do not know.

10          Q       All right. Lbu. I thought it was internal  
11 operations, but I wasn't sure. So any --

12          A       I don't know.

13          Q       Oh, okay.

14                   Ibubiz.

15          A       Yeah. I don't know how to pronounce that  
16 either.

17          Q       All right. Again, I thought it was, again,  
18 an internal tool since it says "always be allocated  
19 in different racks."

20                   Is that a physical rack, like server rack?

21          A       Yeah. That would be my understanding.

22          Q       Okay. So it's likely -- most likely, it's  
23 an internal tool.

24                   Leaf4 and particularly the part I'm  
25 interested in is stores action metadata on actor

1 ID.

2 I didn't understand exactly how user data  
3 is used here or if it's made accessible or shared  
4 with third parties with this.

5 A I don't know. I'm not familiar with the  
6 Leaf data system.

7 Q Okay. Leaf5 seems like it's an internal  
8 tool or no?

9 A Yeah. I don't -- I don't know about Leaf5  
10 either, other than what's written there.

11 Q When it says in-memory cache for actions  
12 and objects, though, that is -- it says it doesn't  
13 actually store any physical objects such as photos  
14 or posts.

15 So I wasn't sure exactly what user data --  
16 is user data just temporarily stored in that system?

17 A Yeah. I'm not sure.

18 Q Okay. Legal SRT Content Takedown.

19 A Yeah. I'm not sure.

20 So I have familiarity with the function of  
21 this system and, you know, as well as what's  
22 described there, but I don't know beyond that.

23 Q But it's not likely; right? I assume it  
24 wouldn't be likely that Facebook would be making the  
25 system -- sharing any of the data in there or making

1 it accessible to third parties, but --

2 A I think that's a fair -- you know, a fair  
3 assumption, but I don't know. I can't confirm.

4 Q All right. Live streaming data or, sorry,  
5 LIV3.

6 A I'm not familiar with that system either.

7 Q All right. Okay. Live Video Delivery, is  
8 that accessible to third parties or shared or would  
9 you know?

10 A I don't know beyond what's written in that  
11 description.

12 Q Okay. And Live Video Processing, is that  
13 user video? Can you just explain to me how user  
14 data fits into that one and Live Video Delivery?

15 A Yeah. So live video would be our -- kind  
16 of our -- a product feature to be able to stream  
17 videos through our products. And as described here,  
18 the Live Video Processing, I can't -- you know,  
19 aside from what's written here, I can't confirm  
20 whether or not there's any external third-party  
21 access to this.

22 Q Or if it's shared or made accessible. All  
23 right.

24 A I'm not sure.

25 Q What about for Lidbpostmortem Cache? It

1 says "internally at Facebook." I assume that means  
2 it's not available or made accessible to third  
3 parties, but can you confirm that?

4 A I'm not familiar with this system. So I  
5 can't confidently confirm.

6 Q This system [REDACTED] how is user data  
7 again captured in here?

8 A So, you know, by that description of  
9 unstructured logging data generated by all  
10 applications across Facebook, what I understand  
11 there is that it is within the application's  
12 determination of what gets logged, which could be  
13 user data.

14 Q When you say "application," is that  
15 third-party applications or Facebook applications?

16 A Let's see what it says there. Yeah. I'm  
17 not -- I'm not familiar enough with it to  
18 definitively say.

19 Q And then would this service be available or  
20 made available or accessible off -- you know, or  
21 made discoverable off Facebook's platform, like  
22 meaning to third parties?

23 A I would find that highly unlikely, but I  
24 can't definitively say.

25 Q Okay. LogDevice, I interpret that to be an

1 internal-only tool for managing operations. Is that  
2 accurate?

3 A That's accurate. I am familiar with

4 [REDACTED]

5 Q And [REDACTED], same?  
6 I wasn't sure how user data got involved in this.  
7 So that's why I'm asking.

8 A Yeah. I don't -- I don't have -- I don't  
9 have familiarity with that.

10 Q Okay. All right. I don't have a clock. I  
11 don't know how long we've been going. I apologize.

12 A It's 1:26 West Coast time.

13 Q We could finish this page, and if you want  
14 to take a break for a few minutes and use the  
15 facilities, the parties can go to their respective  
16 rooms when we finish this, and I think we'll  
17 probably, pending no surprises, finish in hopefully  
18 an hour or less.

19 Does that work, Mr. Pope?

20 A That's fine.

21 Q If you have time today, I'd love to get it  
22 all done today, if possible.

23 A Sure. Yeah. I'm happy to work with you on  
24 that.

25 Q All right. And then Manifold.



1 A So this is not --

2 Well, maybe if you can ask the question  
3 again.

4 Q So Manifold, the reason why I'm asking is  
5 it says it's [REDACTED],

6 [REDACTED] [REDACTED]  
[REDACTED]

8 A Yes.

9 Q It is analogous to Amazon S3.

10 A That's correct.

11 Q I didn't understand, A, how user data was  
12 stored in it; and if it's accessible -- you know, if  
13 you know if Facebook either makes it accessible or  
14 shares that user data in that system.

15 A Yeah. You can -- you can think of this the  
16 same as EverStore. It runs on top of [REDACTED] and  
17 [REDACTED] where it provides additional, like, S3-type  
18 features to the data stored that -- [REDACTED]

[REDACTED] [REDACTED]  
[REDACTED]

21 Q Has Ever- -- I got it.

22 And just to be clear, EverStore is --  
23 Facebook doesn't, you know, share or make the data  
24 in EverStore accessible to third parties or off  
25 platform or does it?

1           A     It does not. It's similar, it's the same  
2     answer there as EverStore.

3           Q     Got it.

4                     And then Memcache.

5           A     Yeah. That would be used internal only to  
6     accelerate performance of applications and products.

7           Q     Facebook applications, not third-party  
8     applications; correct?

9           A     That's right.

10                  MR. GARRIE: Okay. We'll take -- we'll go  
11     off the record.

12                     (Recess taken.)

13                     (Off the record at 1:29 p.m. Back on the  
14     record at 1:46 p.m.)

15                  MR. GARRIE: Okay.

16           Q     Metastore, managing metadata for the Data  
17     Warehouse.

18                     Any idea here, Mr. Pope?

19           A     No. I'm not familiar with that system.

20           Q     Okay. And Mononoke?

21           A     Mononoke, yeah. Yeah. I'm not -- I just  
22     know that this is basically what's written there,  
23     version control server.

24           Q     How would user data be involved? Just to  
25     understand it for myself, like, any idea how user

1 data is involved in that?

2 A Yeah. That would be the -- this is  
3 actually comparable to the Fbcode one we covered  
4 earlier as being a version control system. So it  
5 wouldn't be a common case; but it's possible to, you  
6 know, submit maybe like a test, the unit test that  
7 would have some user data. That would be the case.

8 Q I get it.

9 When you say "submit," it would be from  
10 Facebook developers, not third-party developers?

11 A That's right.

12 Q Multi Party Live. Any ideas, Mr. Pope?

13 A No. I'm not familiar with this system  
14 either.

15 Q I'm fairly familiar with MySQL, but --

16 A Uh-huh.

17 Q -- is it accessible to third parties? I  
18 mean, it can be.

19 A So this is the general category that  
20 covers, I believe, multiple instances of the  
21 database and I'm not familiar with the specifics.

22 Q Got it.

23 New Product Experiments. I just assume --  
24 well, any -- is this internal new product? Like it  
25 says "small entrepreneurial teams." Is that

1 Facebook teams?

2 A Right. Yeah. This is the NPE referenced  
3 earlier. This is kind of the team-level view of  
4 that where they have some storage provisioned for  
5 that team or for these new product teams.

6 Q But it's Facebook internal teams that are  
7 accessing data, not third-party developers that are  
8 building new product experiments and accessing the  
9 data there?

10 A Yeah. I don't actually know that level. I  
11 can't really confirm that.

12 Q Okay. And Novi Gateway Transaction  
13 Storage.

14 A Yeah. Aside from what's written there is  
15 what I know.

16 Q Okay. So we don't know if that data in  
17 there is being -- because it says a queueing system  
18 storing user data.

19 So you don't know if Facebook is making  
20 that accessible or sharing that with others outside  
21 of Facebook?

22 A That's correct. I don't know how it's  
23 used.

24 MR. GARRIE: Okay. So, Facebook, can you  
25 send the letter to the court reporter so she has it

1 to enter as an Exhibit A to the proceeding, just the  
2 letter with the exhibit rather than the entire  
3 declaration?

4 MS. CLARK: You're just saying the  
5 Exhibit A to the declaration, or are you asking for  
6 the separate file name that had the details of the  
7 system?

8 MR. GARRIE: The spreadsheet. Ideally, the  
9 spreadsheet so she can have the spellings of the  
10 different terms. The document we're looking at  
11 right now.

12 MS. CLARK: Okay. Absolutely.

13 MR. GARRIE: All right.

14 MS. CLARK: Is her email address in the  
15 chat?

16 THE COURT REPORTER: Yes. And also,  
17 counsel on the line, if you can go ahead and put  
18 your email addresses in the chat so I can follow up  
19 later to find out your appearance information.

20 MS. CLARK: And we're sending this to you  
21 right now.

22 MR. GARRIE: So we'll go back on the  
23 record.

24 Q The Oculus Recommendation Events, I assume  
25 that's the Oculus system and that's not the Facebook

1 platform or are they combined?

2 A I don't know the answer to that in terms of  
3 the last part of being combined. But this is  
4 referring to Oculus, the reality labs product.

5 Q Okay. And OFFD. This is kind of a very  
6 generic description, storing off-line messages.

7 Is that for users, employees? Who can  
8 access it? Do you have any idea, Mr. Pope?

9 A No. I'm not familiar with that system.

10 Q Okay. OIL. Any idea, Mr. Pope?

11 A I'm familiar with the service, but I don't  
12 actually know really more than what's written there.

13 Q Is the service available to third parties?

14 A I don't know.

15 Q And the reason I ask, it says it provides  
16 the user writing the data with the flexibility to  
17 store the bytes on the storage systems of their  
18 choice, while solving the reader's "where do I get  
19 that data" problem.

20 A Uh-huh.

21 Q So I didn't know if this was a service you  
22 were making available to third-party app developers  
23 or if this was a service that was only available and  
24 accessible -- to where the user's data is only  
25 accessible and shared for Facebook, if that makes

1 sense.

2 A Yeah. No. That makes sense.

3 Yeah. Yeah. I just don't have enough  
4 knowledge of that system to say.

5 Q Okay. Omnistore. Any idea here?

6 A No. I just have a very -- really, the  
7 level of knowledge that's written there is the limit  
8 of what I know about it.

9 Q And the way I read it is that it's only  
10 data storage and sync services to support Facebook's  
11 products.

12 My question is, is the data store also  
13 available or made accessible to third parties? I  
14 couldn't tell from the description.

15 A Yeah. Yeah. I'm not familiar enough with  
16 it to know, to be honest.

17 Q And the reason I ask is it says "messenger  
18 contacts." So that could obviously go one way or  
19 the other.

20 A Uh-huh. Uh-huh. Yeah. I just haven't  
21 worked with it to know.

22 Q That's all right. "I don't know" is  
23 perfectly acceptable. I will follow up with  
24 Facebook, as appropriate.

25 Opportunity Store, which sits on [REDACTED]

1 What's Lift?

2 A Yeah. I'm trying to remember if I know  
3 anything about this one. Yeah. I don't recall.  
4 I -- I recognize the name, but this was -- it's been  
5 some time since I worked with it. So I don't recall  
6 the details of it --

7 Q This is certainly --

8 A -- nor do I remember.

9 Q But this does have user data that is in --  
10 the user data relates to, is it, ad impressions and  
11 clicks?

12 A Yeah.

13 Q Conversions?

14 A I'm not sure. I think. Again, yeah, I do  
15 think that that's a reasonable interpretation, but I  
16 can't confirm.

17 Q Got it.

18 But it's not clear if it's an internal tool  
19 that then -- or if it's available to advertisers or  
20 third-party app developers?

21 A Right. Right.

22 Q So you're not sure.

23 Orderdb. What's historical transaction  
24 data?

25 A So I do know that this is related to the



1 payment system. But beyond that I don't know, you  
2 know, more than what's written there. But I can  
3 share that the context would be related to the  
4 payment system.

5 Q And at the time the system was being used,  
6 was the data in it --

7 I guess you may not know the answer, but  
8 I'll ask it anyways.

9 Was the transactional data available or  
10 made accessible to third parties or app developers  
11 or people off the platform, the user transactional  
12 data?

13 A Yeah. I'm not sure on that one. I don't  
14 know.

15 Q And then Payments. That's built on the  
16 MySQL. So the above MySQL, this is what -- a system  
17 that's running on that. Is that how I should read  
18 that?

19 A Yeah. That's actually correct.

20 Q Now, do you know if this is made accessible  
21 or shared, the data in this system, with third  
22 parties, the user data?

23 A Yeah. I don't know.

24 Q Do you know if this went into effect after  
25 2014?

1 A I don't know that one either.

2 Q Perforce. Is that how I say that?

3 A Perforce, yeah.

4 Q Perforce. I don't even know what a  
5 Perforce user is. So maybe you can tell me. What's  
6 a Perforce user?

7 A The Perforce that I think they're referring  
8 to here is another kind of version control system,  
9 and I don't -- but I don't know beyond that.

10 Q But a Perforce user is not a consumer user.  
11 It's a Facebook developer or -- the point is that  
12 the Perforce users are a subset of a Facebook  
13 employee?

14 A Right. Right. And kind of as stated  
15 there, you know, Perforce users throughout the  
16 company. So it's supporting those internal  
17 employees is my best understanding of it.

18 Q All right. And Periplus. How is user data  
19 involved in the cache for the news feed data?

20 A I have -- I have some recollection of  
21 working with this one, but, yeah, aside from just,  
22 you know, inferring that this is the cache for news  
23 feed, news feed would contain, you know, user data  
24 posts and so forth, but I don't know beyond that.

25 Q But is it temporary cache or is it --

1           A     I don't know. Yeah. That would be my  
2     assumption, but I can't say definitively.

3           Q     And then Pinpoint, looks like it's an  
4     internal tool for infrastructure support team. Is  
5     that accurate?

6           A     I'm not familiar with this, this system.

7           Q     Well, "Why did Facebook DAP go down between  
8     today and seven days ago?"

9           A     Yeah.

10          Q     Is there any reason, as far as you can  
11     think of, that a third party or somebody --

12                Is there any reason Facebook would make  
13     this accessible, the user data in this system  
14     accessible to a third party?

15          A     I don't know of any, but I'm not really  
16     qualified to confirm.

17          Q     But when it says "metric movements," what's  
18     that mean? Do you have any idea?

19          A     Yeah. So for that specific example, you  
20     know, why did DAP. So that would be daily active  
21     people. You know, we moved away from using the term  
22     "daily active users," for whatever reason. So why  
23     did the Facebook daily active people go down between  
24     today and seven days ago? So analyzing the usage  
25     patterns of our products is what that's referring

1 to.

2 Q Got it. So all right.

3 So maybe Facebook would make it available  
4 to -- but is it anonymized? Do you know if the  
5 data -- is this -- when we have user data, is this  
6 de-anonymized user data or is this anonymized user  
7 data? Just generally when people say, "We have user  
8 data," is it specific individual users, like  
9 Joe Smith?

10 A Yes. We don't actually have that, yeah,  
11 level of. Like that wasn't a part of kind of when  
12 we asked what type of data is stored there. I think  
13 you'll have a great deal of variance across the use  
14 cases.

15 So given an application's purpose, it could  
16 very well be storing anonymized data, but we don't  
17 know. Like we don't actually have visibility to  
18 that.

19 Q Well, when you sent out the survey, I  
20 assume you told people what user data is, no?

21 A Uh-huh. Yes. We provided examples and  
22 references for them to confirm.

23 Q Can you provide a copy of the surveys that  
24 were sent out so maybe that will clarify what or how  
25 you defined user data?

1           Because I'm assuming that the user data  
2   isn't anonymized data. This is actual user data and  
3   not anonymized data because you wouldn't have to  
4   delete anonymized data from a system, no?

5           A     Correct. This would be, you know, data we  
6   collect about users using our products like through,  
7   you know, them uploading or other things that we  
8   understand by their use of our products.

9           Q     Yeah. Through the use of apps or whatever  
10   it is. But the point is is that it's user data,  
11   meaning that it's identifiable to a user --

12          A     Right.

13          Q     -- or people or however you --

14          A     Yeah.

15          Q     I just want to make sure. This entire --

16          A     Go ahead. Sorry.

17          Q     No. No. I just wanted to make sure  
18   because when I say "user data," right, I mean  
19   identifiable information to an individual or a  
20   person.

21          A     Uh-huh.

22          Q     So I'm assuming for all of these systems,  
23   they have identifiable user, meaning identifiable to  
24   an individual like person in the system.

25          A     Our definition of user data is a little

1 more broad, I think, in that it might -- you know,  
2 it may be things that we collect by the use of our  
3 products and I can't say necessarily if it meets  
4 your -- that more specific criteria of being  
5 identifiable to a specific person.

6 Q Well, identifiable to a Facebook ID or you  
7 have a universal ID; right? Identifiable at some  
8 level back to an individual.

9 Well, when you send the questionnaire over,  
10 I'll be able to see to make sure we're looking at  
11 apples to apples because it could be, Counsel  
12 Kutscher Clark, just to make sure that this is  
13 actual user data so we're not wasting everybody's  
14 time looking at it.

15 MS. CLARK: So I'm not sure if Mr. Pope  
16 understood the question, but this is what I  
17 understood you to be asking, Mr. Garrie.

18 You're trying to figure out if all of the  
19 data in these sources have called user data would be  
20 data linkable to a specific person or whether it  
21 might also be data that's aggregated with data about  
22 other people or data that's been anonymized so it  
23 can't be linked back. Is that what you're trying to  
24 get at?

25 MR. GARRIE: Yeah. Because when you

1 submitted the declaration, right, it said that for  
2 user data, and I just -- I've been assuming just --  
3 I just thought I'd dot my i's and cross my t's.

4 When I say "user data," I'm talking about  
5 whether there's some entity that has a first name  
6 and last name or a Facebook account. We'll just  
7 leave it at that. Has a Facebook account credential  
8 that is hopefully a human breathing entity, but  
9 whatever.

10 The bottom line is is that when you say  
11 "user data," that that type of data, in some  
12 fashion, exists. So, for example, cache or news  
13 feed data is the cache for a particular user's news  
14 feed. Now, it may only be available internally,  
15 never shared and accessible. I just want to make  
16 sure, because if it's anonymized or genomized or  
17 tokenized or anything like that.

18 So maybe when you send the survey, if you  
19 could just send the -- to counsel the survey and  
20 read it over, but it would be very informative for  
21 me to make sure we're actually looking at user data.

22 THE WITNESS: Yeah. So just to clarify on  
23 that request.

24 So we do have a standard definition. The  
25 actual forms that were sent would be data system

1 specific. So how Pinpoint actually did this versus,  
2 you know, Workforce or what have you may be  
3 different; right? So there's multiple. It's not  
4 like a uniform form, you know, one single form that  
5 was sent to the applications and products using the  
6 data system.

7 I can, you know, share some examples of  
8 user data that might help.

9 One is, you know, data provided by our  
10 users, like user-generated content, posts, so forth;  
11 data we observe about a user, such as location data,  
12 network data; the data we infer about our users',  
13 like, interests. These are all things that we would  
14 consider user data.

15 BY MR. GARRIE:

16 Q But that's all associated with the user?

17 A Right.

18 Q So that can infer data about people  
19 anonymi- -- like the key part of it is that it goes  
20 back to an individual user, an individual discrete  
21 entity?

22 A Right. Right. So where I'm kind of  
23 pausing on this and I'll maybe have to take this  
24 question back to give you a more vetted answer;  
25 right?



1 Q Okay.

2 A Yeah. Is that these will tie to something  
3 specific. I'm not sure that it would be  
4 specifically a user or maybe it's tied to their  
5 photo. But I don't want -- yeah. I don't want to  
6 speculate here kind of the distinction like maybe  
7 what's covered when I answer this question. It  
8 might be a little bit more broad than the way that  
9 you're thinking of specifically.

10 Q I mean, that's what's occurring to me as  
11 we're going through, that may be how you define user  
12 data. So let me think about it.

13 We'll revisit it at the end, but yeah.  
14 Because the reason why I asked is where it says  
15 "root cause analysis for metric movement," that's  
16 what really got me thinking how user data, is that  
17 individual users, you know, active pool and you have  
18 all their IDs in there? I mean, how you identify  
19 them is different, right, whether you use a hash, a  
20 UID or a string of numbers, letters, gibberish.  
21 Whatever, however it is, right, the key point being  
22 is that it's a user; right? And that's why --  
23 that's really why I asked the follow-on question,  
24 but we can revisit it at the end because we're  
25 close, and the finish line is there. So we'll

1 revisit it in a second.

2 A Yeah. Yeah. But just in terms of what we  
3 provide to you, I'm not sure what that process is in  
4 terms of making a request and going through whatever  
5 the legal channels are. So I don't want to misstep  
6 or go out of protocol.

7 Q No. So let me help you out.

8 Facebook has an internal process working  
9 with Ms. Kutscher Clark, who is very familiar,  
10 fantastic outside counsel. She will help you and  
11 your client. I believe you actually have in-house  
12 counsel still on. There's a lot of people. But  
13 Counsel Chen is also very competent counsel. They  
14 will advise you exactly on the process.

15 You're just -- I'm asking you questions.  
16 If you don't know the answer, say you don't know.  
17 If you need to get back to me with information, say,  
18 "I want to confer with counsel," and we'll get back  
19 to you with clarity of an answer. There's nothing  
20 wrong with that.

21 A Yeah. Sounds good.

22 Q The reason why I'm asking is I just want to  
23 make sure we're talking apples to apples for the  
24 follow-on.

25 Poseidon. Poseidon sits on top of [REDACTED]

1 and it's designed to store user side features.

2 Do you have any idea if that's --

3 A No.

4 Q -- accessible or made available?

5 A Yeah. I'm not -- I'm not familiar enough  
6 with that system.

7 Q Because an API could have access to a  
8 storage service or third-party app developer or --  
9 or not. So okay.

10 A Yeah. I'm not --

11 Q Presence Storage Service, green dot. I  
12 didn't know what that meant.

13 A Yeah. I'm not -- I'm not sure what is  
14 meant by "green dot" either. I think it's --

15 Q Do you know what a power user is? Is that  
16 an internal user?

17 A I don't know. I'm not sure what they mean  
18 by that.

19 Q Puma, I assume this is only available  
20 internally or is it made available or accessible or  
21 shared with third parties?

22 A Yeah. I'm not -- I'm not -- I'm not sure  
23 on that one either.

24 Q Okay. Do you know why they capitalize  
25 Stream Processing?

1           A       I don't know, actually. Yeah. I'm not  
2       sure.

3           Q       Okay. But fully managed general purpose  
4       Stream Processing, I couldn't figure out what that  
5       meant. For me to figure it out --

6           A       So --

7           Q       I mean, any thoughts, ideas?

8           A       No. I can't -- yeah. I'd just be  
9       speculating. I don't have contextual knowledge on  
10      why they phrase it that way.

11          Q       Okay. PYMK Leaf Tailer. What's a leaf  
12      machine?

13          A       Yeah. I'm not familiar with this  
14      technology either or the leaf machines.

15          Q       But you are sure that it has user data in  
16      it somehow?

17          A       That it could have user data.

18          Q       That it could have use data?

19          A       Yeah. So that's, again, kind of going back  
20      to where you pointed out something that should have  
21      been in that declaration is that the questions asked  
22      was, is it possible to store, you know, or something  
23      along these lines, is it possible that an  
24      application of products using your data system could  
25      store user data?

1 Q Okay. Fair enough.

2 RAlligator. So I get what this is doing.

3 My question is, is this available to third  
4 parties or accessible or made available?

5 A Yeah. I'm not sure about this system. I  
6 don't know the answer to your question for  
7 RAlligator.

8 Q And, again, but you have a point of contact  
9 in your effort through the year that this is -- you  
10 know, or even before we even get to that, you have  
11 some response from them that would explain at some  
12 level more detail what RAlligator is?

13 A Yeah. And just to kinda clarify that path  
14 and why this, you know, has been a pretty cumbersome  
15 effort is we do have the initial point of contact  
16 that support the data systems themselves; but the  
17 products and applications and infrastructure systems  
18 that use those data systems, each one of them have  
19 their knowledge of what data. You know, the actual  
20 point of contact for RAlligator might not know what  
21 data is actually stored there. You would have to  
22 then go to the next order of contacting the  
23 application the product's using.

24 Q Yeah. But they would be able to provide --  
25 but they can provide a general idea. Your point of

1 contact is responsible for the system at some level,  
2 no?

3 A They're responsible for the storage  
4 capabilities and the deletion configuration. So  
5 they would have --

6 Q Are they technical?

7 A What do you mean by that?

8 Q Well, you say they're responsible for it.  
9 So they would know what data is being written to  
10 RAlligator? I mean, I assume you had to ask the  
11 person that actually runs the data system at some  
12 level. Whether or not they know what's put in there  
13 is different, but they would need to know what it  
14 does and how it works, no? To answer your question.

15 A Yeah. Yeah. No. I think so.

16 What they will know, and this is where, you  
17 know, the solution of how they do this will vary  
18 across the different data systems based on the  
19 technologies.

20 So I don't know RAlligator or RAlligator to  
21 be able to say how they do it, if they have that  
22 knowledge. But for a data system that provides  
23 storage services like Zippy, key value store, the  
24 customer's external, but the actual applications and  
25 product that store data in Zippy will be an

1 admission control process that asks is it user data  
2 and they'll check -- and if they check yes, they  
3 have to select the configuration options to delete  
4 the data according to our policies.

5 That's the extent that they know. So they  
6 should be able to --

7 Q Yeah. But there's more than -- yeah. But  
8 they'll know if there's user data in there and what  
9 it looks like. Whether it's accessible or not is  
10 downstream from that; right?

11 A My --

12 Q Yeah. That's downstream.

13 First I need to figure out if the user data  
14 is stored in it and then if it's available only  
15 to -- the first step in the analysis is, is this  
16 data, the user data, accessible or shared with third  
17 parties or off Facebook's internal platform or is it  
18 aggregated and shared with another Facebook  
19 application that then is shared with third parties;  
20 right?

21 So, but the first step is just figuring out  
22 like -- like the call logs, if I'm building a  
23 third-party app and I'm making function call logs --  
24 I'm making function calls and you're restoring  
25 because the aggregation of call logs is ambiguous.

1 Like is it aggregating all of a particular call in a  
2 particular function going back to 2007 or is it, you  
3 know, understanding -- and then --

4 But the first simple question is, does  
5 it -- is it available to third parties or accessible  
6 to the outside, and then going from there?

7 A Right. Okay.

8 Q Readstate. When you say "the feed," is  
9 that the user feed, the friend feed? What's a feed?

10 A What's a feed? All the stories user has  
11 seen.

12 So it seems to be referring to stories  
13 through the news feed.

14 Q And do you know if this is -- is this  
15 available to third parties or off-platform or shared  
16 or accessible?

17 A I don't know. Yeah. I'm not sure about  
18 that.

19 Q Ready At Dawn I looked up. It's a game  
20 studio.

21 A Uh-huh.

22 Q What user data is this -- is this Facebook?  
23 I mean, can you clarify when you say a --

24 A Yeah. I do have some familiarity with  
25 Ready at Dawn.



1           So it was an acquired company, a gaming  
2       studio for our reality labs division, and this would  
3       be data -- the data that we talked about there was  
4       the users of that game. Like so you log in and, you  
5       know, you have some kind of profile.

6           Q     Do they log in via Facebook or are they  
7       logging in via -- you know, how are they when we say  
8       user?

9           A     Yeah. I don't -- I don't actually know  
10      that detail.

11          Q     More importantly, is that data, the user  
12      data, available to any third parties?

13          A     Yeah. And I'm not -- I'm not familiar  
14      enough with that detail of the game and that system.  
15      Yeah.

16          Q     It's a video game that -- it's the user  
17      data about the people that are playing the game,  
18      though?

19          A     Right.

20          Q     Recon DB.

21          A     Yeah. I just know, you know, kind of  
22      what's written there is my knowledge of the system.

23          Q     Would it be likely that Facebook would give  
24      third parties access to their payment reconciliation  
25      for users?

1           A     Yeah. That's outside of my kind of  
2     experience or expertise to actually have an opinion  
3     on that.

4           Q     But when you emailed it, when you sent them  
5     the questionnaire, you asked does it have user data,  
6     and they can tell you if it does and how it gets  
7     deleted?

8           A     Right.

9           Q     Reliable Volumes.

10          A     I do have some knowledge of this system.  
11     Yeah. I do have some knowledge of this. So this is  
12     more like internal storage for other systems.

13          Q     Yeah. I assume you're not going to let  
14     Google access --

15          A     Right.

16          Q     -- this.

17          A     So this operates at a lower level in the  
18     tech stack, you know, providing more storage  
19     capabilities, like physical storage.

20          Q     Yeah. All right. All right. So then  
21     Research Home?

22          A     Uh-huh.

23          Q     Who are UX researchers? Are they Facebook  
24     researchers or are they third-party people?

25          A     Yeah. I'm not sure who those people are.

1 Q And they're storing user data in the  
2 system?

3 A Yeah. So I do have some experience with  
4 just the Research Home in terms of that. This would  
5 be for user experience research where they're maybe  
6 doing a case -- or, you know, came in and -- you  
7 kinda come in and use the product and what have you.  
8 That might be one example of the type of research  
9 that's being done.

10 Q All right. But is it being done by and  
11 made available to third-party developers or anybody  
12 else?

13 A That I don't know.

14 Q Okay. RocksDB. It says anyone. Does that  
15 mean anybody can access this service?

16 A So this is similar to Reliable Volumes. It  
17 is a lower-level storage interface. So it would be  
18 used by other systems.

19 Example, Zippy uses [REDACTED] underneath. So  
20 it's a lower layer in the stack.

21 Q And when it says "anyone," who's anyone?

22 A Employees of Meta.

23 Q Okay.

24 A So they would use it. And just to kinda  
25 clarify that, how they use it, these are libraries

1 that they would actually write their applications  
2 and compile this code to store the data onto some,  
3 you know, storage system.

4 Q So, but third parties wouldn't have access  
5 to that?

6 A No.

7 Q RSC Storage.

8 A Yeah. I mean, the familiarity is just  
9 what's written there that I have with it.

10 Q Scribe. Any --

11 A I don't have any additional knowledge or  
12 experience.

13 Q Scuba?

14 A I do have experience with Scuba and it's an  
15 internal tool for doing visual data analysis. Yeah.

16 Q When you say --

17 A Yeah.

18 Q Internal tool being? Sorry. Could you  
19 repeat that?

20 A Yeah. It would be used by Meta employees.  
21 It's not a production service or tool. It's for,  
22 you know, analyzing. They call it kinda deep diving  
23 on data, like some data that you import to the  
24 system to, you know, kinda like -- like maybe things  
25 like Excel with a bunch of extra features; right?

1 Like you put in a bunch of data and you can like cut  
2 and, you know, kinda do different joins and do  
3 different manipulations to understand the data and  
4 do analysis.

5 Q Yeah. Excel on steroids.

6 All right. SecureDumbo. That to me is  
7 pretty obvious.

8 But third parties don't have access, I  
9 assume, to your backups?

10 A Right. Yeah. They do not.

11 Q ServiceFriend.

12 A Yeah. I'm not familiar with that, that  
13 system.

14 Q And ServiceLab, if I built like an app,  
15 would I get access to this?

16 A I actually don't have enough familiarity to  
17 say on this one.

18 Q Okay. SFTP, I assume you don't provide  
19 access to that or do you? It says within the  
20 company and with external vendors.

21 A Yeah. So it is the method that's used for  
22 exchanging files with other external vendors --

23 Q With --

24 A -- of the --

25 Q With user data?

1           A       There could be user data, you know, they  
2       indicated by saying this on the form. I don't know  
3       the nature of what those use cases are, though,  
4       where they would be doing that.

5           Q       Got it.

6                   Is SFTP a service or something they back up  
7       or -- well, we can figure it out later.

8                   ShadowDB.

9           A       Yeah. So this is providing kind of a test  
10       MySQL instance.

11          Q       It says it captures all incoming traffic  
12       and replays it on another test in real time.

13                   Who is this available to or accessible to;  
14       right? Because this is user -- if I'm assuming it's  
15       user data, how is user data -- is this user incoming  
16       traffic of how they're using the Facebook platform?

17          A       I'm not actually sure. Yeah. No. I don't  
18       know.

19          Q       Okay. I assume Sisyphus is not, but can  
20       you confirm that it's not accessible to third  
21       parties?

22          A       Yeah. My experience with Sisyphus is that  
23       it's data processing jobs that were a message DB.  
24       So this is part of a messaging service and does kind  
25       of batch data processing, which would be an internal

1 process.

2 Q All right. Novi, that's your crypto;  
3 right?

4 A Correct.

5 Q Do you know if that's accessible to third  
6 parties?

7 A I do not know.

8 Q Okay. Sparse Inner Product. I didn't know  
9 what candidate generation meant.

10 A Yeah. I'm not -- I'm not familiar with  
11 that system or to be able to, yeah, really clarify  
12 that term.

13 Q Sparta, this is an application developer  
14 review tool. Is that to who?

15 A Yeah. I'm not familiar with this one  
16 either.

17 Q Okay. All right. So you're not familiar  
18 if Facebook allows other third parties or shares? I  
19 mean, I would assume it doesn't, but okay.

20 A Yeah. I'm not sure.

21 Q All right. Stylus.

22 A Yeah. I am familiar with Stylus, I mean,  
23 to kind of the extent that they're a stream  
24 processing library. So it would be used by internal  
25 software developers, internal being Meta employees,

1 to create the stream processing jobs.

2 Q Got it.

3 Swift. I understood that to be an internal  
4 Facebook dev tool or ops tool, but can you confirm  
5 that?

6 A By the description with EE Storage, you  
7 know, these are for our company use by Meta  
8 employees and as well that this is for archival  
9 workloads, such as backups and data warehousing.

10 I don't have direct experience with Swift,  
11 but I think by this description is it's an internal  
12 system.

13 Q And then I would -- yeah. I would agree.

14 Synapse, I assume is similar, but again --

15 A Yeah. I'm not familiar with this system to  
16 be able to answer.

17 Q What's composability mean?

18 A I'm not sure. I'm not familiar with that  
19 term in this context.

20 Q Tableau I've worked with, but do you give  
21 third parties access to your Tableau?

22 A I don't know, actually. I'm not sure.

23 Q But user data is moved into Tableau and  
24 then you run BI on it?

25 A There can, yeah. So there can be user data



1 in performing the Tableau feature analysis.

2 Q Yeah. Okay. Always a selling point.

3 A Yeah.

4 Q TacoSRS. Any idea?

5 A No. No. I'm not sure. I don't know on  
6 that one.

7 Q Is there any reason you would think that  
8 Facebook would make this accessible to a third  
9 party or the user data stored in this system or  
10 shared?

11 A Yeah. I think I'm just actually not  
12 familiar enough with its use to be qualified to  
13 really --

14 Q Because it says --

15 A I could guess.

16 Q And the reason I ask is it says "now  
17 deprecated Taco created as part of an effort."

18 So I took deprecated to mean no longer in  
19 use, but maybe I -- well, fair enough.

20 Tally.

21 A Yeah. I'm familiar with the service, but  
22 can't really speak to its use cases or...

23 Q How? What's your familiarity with the  
24 service?

25 A It's a service that runs on top of [REDACTED]

1 and basically is an application you can configure to  
2 use for storing counters. Like an example I think  
3 was given while we did talk with Tally would be the  
4 number of likes that a post has received.

5 Q All right. So then would you know if  
6 third-party developers were just --

7 Do you know if Facebook makes this  
8 accessible or shares it off Facebook's platform,  
9 access to this data?

10 A Yeah. I don't know. I don't know that  
11 level of detail.

12 Q All right. I assume that's the same for  
13 Tally on Sigma. That's the same --

14 A Yeah.

15 Q -- or is it different?

16 A Yeah. I don't remember the distinction  
17 between these two. They're related, but I don't  
18 remember how they differ.

19 Q TAO?

20 A Uh-huh.

21 Q I get what it is.

22 My question is, would you know if Facebook  
23 makes it available or accessible or shares this data  
24 in this -- the user data in TAO with any  
25 off-Facebook platform?

1           A       Yeah. I don't have, like, familiarity with  
2       that and how that all ties together with the social  
3       graph and to the question that you're asking. So I  
4       don't know how that...

5           Q       Threat Profile. It says it's an internal  
6       tool, so I just assume it's not shared.

7                   Is that an accurate assumption?

8           A       Yeah. I -- I -- again, I'm not -- you  
9       know, this isn't my area of work, but I think from  
10      the description --

11          Q       All right.

12          A       -- that's pretty clear. Yeah.

13          Q       All right. TimelineDB. I get what it is.  
14                  Would you, by chance, know if it's shared  
15      or made accessible to --

16          A       No.

17          Q       -- off --

18          A       Yeah. No. I don't know.

19          Q       And then TRU. I assume it's not.

20          A       Yeah. Yeah. You're right. So this is  
21      more internally used to understand our -- our  
22      products and systems.

23          Q       And same with TW Storage?

24          A       So -- right. So just to make sure we're on  
25      the same, yeah, so this would be where --

1 Tupperware is our container-based service. So this  
2 would be storage tied to those containers. So it's  
3 more of a lower-level internal --

4 Q Yeah.

5 A -- technology. Yeah.

6 Q Yeah. I would be surprised if you made  
7 that accessible.

8 And same with UDB?

9 A Yeah. This is the backing, kind of, as it  
10 says, backing database for TAO, [REDACTED]

[REDACTED]

[REDACTED]

13 I don't know. Again, kind of like I don't  
14 know the linking in terms of availability to the  
15 social graph, but that's like the underlying, you  
16 know --

17 Q And the social graph --

18 A -- storage.

19 Q And the social graph data sits in the TAO;  
20 right?

21 A Right. So that's like kind of our primary  
22 cache for driving our applications and a large part  
23 of that is where the social graphing is.

24 But again, you know, I don't have like  
25 the -- that's definitely not my area of expertise

1 getting to how those all tie together.

2 Q Yeah. And so where it says "used for  
3 testing," is that internal testing for UnittestDB?

4 A That is actually my understanding of it and  
5 I did have some work with UnittestDB, but my memory  
6 is a bit dated to be truly definitive on that.

7 Q All right. Your notes would reflect it at  
8 some level?

9 A Yeah.

10 Q And Up2x?

11 A Yeah. I'm not -- I'm not too familiar with  
12 Up2x.

13 Q All right. What about UIP?

14 A Yeah. I'm not familiar with UIP.

15 Q Can you explain to me, what is human  
16 readable? I hear the term a lot, human readable.  
17 To me data is data I can read, binary or hexadecimal  
18 or words.

19 A Yeah.

20 Q When they say "human readable," what does  
21 that mean?

22 A I think that that's like written in a  
23 common human language, you know, rather than, you  
24 know, kind of encoded, you know.

25 Q But someone translates the coding to

1 understand it; right? I mean, machine readable  
2 means a machine can translate the data, right, or am  
3 I reading that wrong?

4 A I can't --

5 Q It's all -- well, human readable is still  
6 data, right, and machine readable is still data?  
7 There's no -- it's still data we're talking about;  
8 right? There's no distinction there.

9 A Right. So I can't really speak to what  
10 whoever wrote -- provided this description is trying  
11 to convey with the human readable understanding  
12 piece.

13 But to your point, I think more relevant is  
14 that, yeah, for our exercise that we went through  
15 and for the data we gathered, this is talking about  
16 data. And, you know, whether it's, you know, in  
17 sentences that we -- [REDACTED]  
18 some kind of, you know, [REDACTED]

[REDACTED] isn't  
20 really relevant to the work that we did. We're  
21 just, you know --

22 Q Right. That's why I was asking.

23 A -- gathering data. Yeah.

24 Q This is the only place for user data that  
25 it says "human readable." So I didn't know if there

1 was a distinction there.

2 A Yeah. I don't think any meaningful  
3 distinction. At least not in the context of the  
4 work that we were doing. We were just concerned is  
5 there data? True/false.

6 Q Got it.

7 And then you wouldn't --

8 Did they clarify if this human readable of  
9 user interests to improve personalization in various  
10 user-facing products, is that just Facebook's  
11 products or is this available to all app developers?

12 A Yeah. I don't -- I don't actually know  
13 about at that.

14 Q Wouldn't they differentiate or identify if  
15 it was non-Facebook products or no?

16 A Yeah. I'd probably be speculating to say.

17 Q The people that wrote it and said  
18 "user-facing products," they're talking about the  
19 products that -- I'm just trying to understand when  
20 they say "user-facing products" what the person, the  
21 context they're responding from.

22 A Yeah.

23 Q You sent them a questionnaire and they  
24 wrote you back, "user-facing products." Is that  
25 Facebook users or you're just not sure?

1 A I'm just not sure and I didn't --

2 Q Okay.

3 A I don't think --

4 Q I read --

5 A So yeah.

6 Q And then a tax application?

7 A I'm not familiar with this product or this  
8 service or system.

9 Q It's calculating taxes, right, for your --

10 A Yeah.

11 Q So Warm Storage, my question for this is do  
12 you make Warm Storage available off -- the user data  
13 in Warm Storage available to third -- off Facebook,  
14 you know, accessible to or off Facebook platform?

15 A So yeah. Well, so Warm Storage is, again,  
16 a lower-level storage solution that's used by  
17 systems on top of it. So I can't really say, you  
18 know. I don't know all the specific cases where  
19 Warm Storage is being used as the underlying  
20 storage. It would be through --

21 Q But there's no way --

22 A Yeah. It would be through a data system  
23 running on top of Warm Storage.

24 Q That's my point.

25 There's no way to access Warm Storage. The



1 data -- the user data in Warm Storage still requires  
2 some database app- -- some application to access the  
3 user data that's stored there.

4 Is that not an accurate assumption?

5 A Exactly. Yeah. No. That's a very  
6 well-phrased way to put it. Yeah. That's correct.

7 Q Makes more sense. All right.

8 WhatsApp Deletion Framework.

9 A Uh-huh. Yeah. This Deld, that one that we  
10 previously talked about with respect to WhatsApp  
11 deletion, the way that deletions happen with  
12 WhatsApp, this is that entire framework of the  
13 component parts. So Deld would fit --

14 Q Yeah. But --

15 A -- into this.

16 Q But that's behind WhatsApp. So a  
17 third-party app. Facebook's not making that  
18 accessible to anybody?

19 A That's correct. That's correct.

20 Q Okay. Wormhole. It says it's limited to  
21 different groups at Facebook.

22 I just want to make sure there's no other  
23 groups that have access to that.

24 A Yeah. I'm not -- I don't have enough  
25 knowledge of it to say.

1 Q On its face, at least it seems to be.

2 A Yeah.

3 Q Okay. XDB.

4 A Uh-huh.

5 Q So I get what it is.

6 My question is, does Facebook make this --  
7 it's, again, a very broad description.

8 Who has access to this self-service  
9 relational database?

10 A Somewhat to the previous statement and the  
11 way you phrased it, this is applications product  
12 developed with provision, a MySQL database through  
13 this XDB service, and there would be an application  
14 running on top of, you know, using tables that are  
15 in an XDB instance. How that application manages  
16 that data is not known at this layer.

17 Q So my question is, for a user -- so the  
18 user data, when they sent this back to you, how did  
19 you figure out what was running on XDB to actually  
20 do something with it?

21 Because I can't imagine you query the  
22 entire XDB platform or the Warm Storage platform and  
23 say delete this one user's data across 20, you know,  
24 however many exobytes or whatever you have of data.

25 A Yeah. So the way it works with XDB, it's

1 self-service. So if I was an application developer  
2 and I wanted to store, you know, using relational  
3 database, I would go and I'd say I want to set up a  
4 table or like, you know, provision some tables for  
5 me, and it would be through that [REDACTED]

6 [REDACTED]  
7 In which case then, okay, [REDACTED]

8 [REDACTED] Like how will those be -- meet  
9 the requirements of the user data deletion policy?

10 Q Then my question is, how many  
11 applications --

12 So but how, then, would they delete the  
13 user data off XDB today? You can't. It would have  
14 to be by application, no?

15 A So the -- the -- one example would be where  
16 the application only needs that data for, you know,  
17 a short amount of time or some determined amount of  
18 time. Let's say five days. And so they can  
19 configure that. Like XDB supports a setting that  
20 says, you know, delete the data in five days or  
21 delete these rows from this table after they're five  
22 days old.

23 I don't know if that kind of speaks to the  
24 question you're asking, but that would be, you know,  
25 one case where through time to live you could --

1 Q And you would assert --

2 A -- delete.

3 Q Facebook doesn't share access to XDB to an  
4 app developer or anyone else that's not a Facebook  
5 employee?

6 A Yeah. Yeah. No. And the way you phrased  
7 that earlier is right. You would have to have some  
8 kind of intermediary application interface to that  
9 data.

10 Q Right.

11 And does Facebook allow third-party app  
12 developers access to the XDB database, data  
13 warehouse?

14 A No.

15 Q Because XDB is a database, no?

16 A Yeah. Basically you can -- you know, if  
17 you want to run a MySQL database, it's kind of a  
18 service that provisions those databases for you or  
19 those instances of a MySQL database.

20 Q All right.

21 A So yeah.

22 Q So XStream. It says "next generation."  
23 Does that mean it hasn't been deployed or --

24 A It is a -- from my experience with XStream,  
25 yeah, it was in early phase of adoption, which would

1 have been, I guess, a year or so ago and, as  
2 I recall, it was, you know, kind of the -- an  
3 improved version, kind of in the space of Stylus,  
4 like when we talked about those libraries where you  
5 create pipelines.

6 Q And XWF-L?

7 A Yeah. I'm not -- I'm not familiar with  
8 this other than XWF is for -- I think it stands for  
9 Wi-Fi, but I don't -- I don't actually have any more  
10 knowledge of that system.

11 Q But it's not a Facebook -- it's not -- by  
12 using Facebook web, you know, social network, I'm  
13 not using X, the Wi-Fi, am I?

14 A Yeah. Unfortunately, I don't know.

15 Q Zippy I get.

16 A Yeah.

17 Q And my only question is, is Zippy  
18 accessible or available or shared with third  
19 parties, the data instances from Zippy?

20 A I don't actually know in that space.

21 Q But it does have user data, right --

22 A It does. It definitely --

23 Q -- when it's aggregated?

24 A Yeah.

25 Q And that's where it aggregates inferences

1 and everything, I guess?

2 A So yeah. Zippy is one of those that we had  
3 previously called in our efforts general purpose  
4 data. So where it's kind of like XDB, if you want a  
5 relational database, go to XDB. If you want a key  
6 value store, you go to Zippy and you go through a  
7 form and it's really serving very many, many  
8 applications. Probably thousands of different  
9 products and services are running on top of Zippy  
10 and using it. Again, it's just a very easy kind of  
11 self-serve key value store technology.

12 Q So then user data that's stored in Zippy  
13 would be -- whether it's accessible or shared with  
14 third parties would be defined by the app that's  
15 sitting on top of it?

16 A So this is where it gets outside of my  
17 expertise and access to the social graph and because  
18 this would be. You know, some portion of this would  
19 be.

20 Q Well --

21 A Yeah.

22 Q Yeah. Some of it would power the social  
23 graph.

24 The other question is, is there other  
25 things? So the social graph provides all the data

1 and that's the DYI tool. If I spin up another  
2 database instance on Zippy and make it accessible  
3 to -- this doesn't feed into the social graph, but  
4 does it have the ability to make it accessible to  
5 third parties or share the data that's stored in  
6 that instance of the -- you know, the database that  
7 you build that doesn't feed the social graph, but  
8 sits outside of the social graph?

9 A Right. No. It would not be.

10 Like maybe restating, I'll state it like  
11 what I understand.

12 If you set up an instance, you get a  
13 provision some portion of ZippyDB and external,  
14 outside of Meta's infrastructure, would not be able  
15 to access the system to use Zippy directly. There  
16 would have to be an internal Meta -- Meta system  
17 that provides an interface to that data that serves  
18 Zippy.

19 Q Yeah. So there would have to be a Meta API  
20 that's not social graph that would give direct  
21 access to that Zippy dataset?

22 A Right. Right.

23 Q Yeah. Makes sense.

24 And who would have access to the API? Who  
25 controls the APIs that grant access like that?

1 A I'm not -- I don't actually know.

2 Q I assume it's one of those code repository  
3 data information. At some level, there has to be  
4 some -- a security has to track the API access  
5 that's granted. Okay.

6 A Yeah. I think that's a fair assumption,  
7 but I don't know.

8 Q Well, they wouldn't let -- I mean, I hope  
9 so.

10 A Yeah.

11 Q Muddler and HulkDB, that's internal  
12 training for the Instagram machine learning or is  
13 that available to third parties?

14 A I'm not familiar with that system to be  
15 able to say.

16 Q Okay. I thought it was alphabetical, but I  
17 was unduly excited.

18 Double Vision. This is Instagram. Again,  
19 is this -- do you know, like, how this system works  
20 and if Facebook makes the data in it accessible?

21 A No. I'm not familiar with Double Vision.

22 Q But it's just Instagram data; right?

23 A That is my understanding from the  
24 description.

25 Q And then IGML SIP?



1           A       Embedding indexing system. Yeah. I  
2       don't -- I don't know anything about this system.

3           Q       AI Compliance Feature Store Features.

4           A       That's a mouthful, but --

5           Q       I didn't know what a privacy-wave task was  
6       and what a user was.

7           A       Yeah. So, yeah, I can clarify those.

8                   So privacy-wave is a privacy-wave task.  
9       Privacy-wave is a framework program that essentially  
10      federates work across the company. So this is  
11      where, for example, XDB, when they want to find out,  
12      you know, who -- you know, if they want to reach all  
13      of the applications and products teams that are  
14      using, they would file those through these waves.  
15      And so it's like kind of a central tooling that you  
16      can get, you know, thousands of tasks, and tasks are  
17      internal ticketing system, you can get thousands of  
18      these tickets centrally managed and then track the  
19      progress. Because it's so highly distributed across  
20      the company, you know, you have hundreds of teams  
21      potentially and thousands of employees that have to  
22      take some action. So that's kind of privacy-wave.

23                   And then the user verification/mediation,  
24      that would be referring to the task owner. So a  
25      Meta employee.

1 Q Okay. Good. That's all I need to know.

2 All right. Then public street-level  
3 imagery and map data platform. I wasn't sure how  
4 user data is involved in that and if third parties  
5 can access the mapping system.

6 A Yeah. I'm not -- I'm not familiar with  
7 that system at all.

8 Q All right. A few more.

9 A Homestretch.

10 Q Fastdesk.

11 A Yeah.

12 Q I assume that's internal customer support?

13 A So that is internal customer support.  
14 However, I don't know if they have any kind of  
15 relationships that would provide, you know, external  
16 access.

17 Q That's fine.

18 RPM, what's that stand for?

19 A I think it's the Red Hat Package  
20 Management, similar to FB Package. It's just a way  
21 of rolling up binaries.

22 Q And when it says it's owned by  
23 Core Systems, what's Core Systems?

24 A That's a division of the company or  
25 organization within the company.

1 MR. GARRIE: Okay. All right. So that was  
2 all of them. I will review the transcript rather  
3 than making you stay on as we go through it all.

4 I thought about it for a minute and what I  
5 was hoping, Mr. Pope, if you could go to a separate  
6 room for a minute or two and I can just speak with  
7 counsel, and then we'll bring you right back in. I  
8 just think that may -- if that's okay. Obviously if  
9 counsel from Facebook wants to go with Mr. Pope or  
10 stay, feel free. But I just want to talk to the  
11 lawyers directly.

12 You did excellent, Mr. Pope. This was  
13 extremely helpful in understanding and I really do  
14 appreciate you taking the time.

15 THE WITNESS: I'm glad I was able to help  
16 you.

17 MR. GARRIE: If you could jump into that  
18 other room before. I just want to talk to the  
19 lawyers about potentially additional questions or  
20 other things; but rather than you hearing them, I'd  
21 rather just understand what may be sought.

22 THE WITNESS: Okay. I'm leaving now.  
23 Thank you. I'll be in the other room.

24 MR. GARRIE: And, Counsel Chen, or anybody  
25 from Facebook, feel free to stay or join.

1           So, Counsel, you can turn your video back  
2     on.

3           So thank you. Counsel or Martie, thank  
4     you. It was actually extremely, very helpful for  
5     me, at least, to get my way through all of it.

6           I will open this with a lot of -- well,  
7     plaintiffs, if there are any other particular  
8     follow-up questions that you think may be  
9     appropriate, I'm willing to hear them. I'm not,  
10    per se, willing to ask them, but I'm certainly  
11    willing to listen.

12           MS. WEAVER: Right. I mean, we're very  
13    appreciative of Mr. Pope's time and understand the  
14    limits of his knowledge.

15           I think in terms of framing for us,  
16    recognizing that not all data is actually  
17    maintained, but there are some records, that is,  
18    there were facilities that collect and made data  
19    available, but those facilities probably are not on  
20    the data source list.

21           And so we think the kinds of questions are,  
22    you know, do these sources have either information,  
23    what is it, is the information in those sources used  
24    in any products that have been made available to  
25    third parties?

1           So some of these questions about is this  
2     data available. Like we know, for example, that  
3     Hive has data in it that has been made available to  
4     third parties, but not because Hive makes it  
5     available.

6           So the simple question of is there data  
7     sitting here in this source and does this source  
8     make it available will miss some obvious  
9     opportunities to learn about what is made available.

10           MR. GARRIE: For Mr. Pope?

11           MS. WEAVER: Yeah. We don't know what  
12     Mr. Pope --

13           MR. GARRIE: Well, Mr. Pope stated up front  
14     that he's not the person with the answer to that  
15     question; right?

16           He said that he asked the simple binary  
17     question, true or false. And data, user data in  
18     this system, like XB or some of these [REDACTED]

19     [REDACTED], I'm assuming of data and  
20     they have applications that, like you said, you can

21     [REDACTED]  
22     [REDACTED] but, as he repeatedly stated, he  
23     doesn't have that knowledge. So it wouldn't be  
24     appropriate to ask him if it is appropriate or not  
25     or so on and so forth.

1 I understand your point and we can talk  
2 about that in a minute, but I do want to let  
3 Mr. Pope go before we go down into that part of the  
4 conversation since I did have a chance to read the  
5 court's transcript on this.

6 So, counsel for Facebook, are there any  
7 questions you'd like to ask Mr. Pope as a follow-up  
8 that I may have forgot to ask? I can't imagine  
9 there is, but maybe.

10 MS. CLARK: No. I don't think there are  
11 any other questions that need to be asked.

12 I will just state for the record that we  
13 request the transcript as marked highly  
14 confidential, attorneys' eyes only, because there  
15 was a lot of very confidential information.

16 MR. GARRIE: Yeah. So noted. And I'll  
17 leave that for the court to work out if the parties  
18 want to challenge any designations and so on and so  
19 forth.

20 MS. WEAVER: We understand that and we  
21 appreciate what has been shared today, and we'll  
22 honor the designation.

23 MR. GARRIE: Okay. So I found it extremely  
24 helpful.

25 But first, why don't we let Mr. Pope go,

1 and then we can resume the second part of the  
2 conversation. Since I have everybody here, maybe I  
3 can streamline some of that as well now that I've  
4 read the court's transcript.

5 So can we bring Mr. Pope back in? Maybe  
6 counsel wants to go and tell him I'm going to bring  
7 him back in. I'm sure he loves lawyers. I just  
8 don't know how much exposure he has to counsel.

9 MS. CLARK: I'll go in.

10 MS. WEAVER: I'm not sure anyone does,  
11 but...

12 MR. GARRIE: Mr. Pope, I wanted to thank  
13 you and let you go enjoy your Friday and your long  
14 weekend, and thank you very much.

15 I may follow up with some additional  
16 questions for particular systems to get the  
17 questionnaires or the data that you got from them.  
18 I, frankly, have to read the transcript to process  
19 everything because we did cover a lot of data. But  
20 I do really appreciate and am grateful for you  
21 making the time, and it was extremely informative  
22 and helpful.

23 So have a good rest of your weekend and  
24 enjoy.

25 THE WITNESS: All right. Well, thank you.

1 Thank you. It was my pleasure, and you're welcome.  
2 Have a great weekend.

3 MR. GARRIE: You, too.

4 THE WITNESS: Okay.

5 MR. GARRIE: Thank you.

6 Okay. So I read the judge's transcript and  
7 I did read the orders. I got lucky that my other  
8 hearing this morning ended two and a half hours  
9 early. I'm not nearly as lucky now.

10 The question I have after reading it and,  
11 Counsel Weaver, to your point, so let me reiterate.  
12 The whole purpose here is to determine a very narrow  
13 issue, right, of determining the nine -- and I'll  
14 read what she said; right? Page 29 of the  
15 transcript says, "Facebook's portion is the only  
16 named plaintiffs. We're only talking about the  
17 named nine now, nine named plaintiffs, which is  
18 fair. That the only data that is discoverable is  
19 that which Facebook intends is shared or made  
20 accessible; and if Facebook says it was not shared  
21 or made accessible, that is the end of the matter."

22 MS. WEAVER: But her order says this;  
23 right? "Apart from the court not having limited  
24 discovery of each named plaintiff's data to only  
25 data Facebook concedes it shared with third parties,



1 the court finds that such a limitation is not  
2 appropriate. As the court previously noted, how  
3 Facebook uses the data it collects about its users  
4 is an open question. The first step in answering  
5 that question is to identify the data it collects  
6 about its users and specifically what it has  
7 collected about the named plaintiffs."

8 So asking Facebook --

9 MR. GARRIE: I was just about to read that  
10 part as well, but thank you for reading it into the  
11 record.

12 MS. WEAVER: Good.

13 MR. GARRIE: So looking at her order and  
14 the additional clarity she did provide, from the  
15 systems we just heard about, I think I -- I'm going  
16 to read the transcript -- maybe 15, maybe 10, maybe  
17 less. I think of them all, there's probably three  
18 or four where that conversation, you know, Facebook  
19 may need to go back and give further clarification  
20 or answers. I could be wrong. I have to actually  
21 read it. We did cover 149 systems. So I'm going to  
22 reserve rights.

23 But I do think that understanding  
24 particularly the ones around -- well, the ones he  
25 couldn't answer, like around the ad networks, around

1 the very large data warehouses where he was -- and  
2 the concern -- and this is all hypothetical. I'll  
3 just be very clear here.

4 It sounds like Facebook has the ability to  
5 answer the questions. I want to do this in an  
6 iterative, reasonable fashion where it's focused but  
7 reasonable, and resolves how the judge framed the  
8 issue.

9 So the first step, from where I'm sitting,  
10 is to go from 149 to a whole lot less; and then from  
11 the whole lot less, get further clarification from  
12 Facebook about what those systems actually do  
13 because it sounds like to me some of these systems  
14 are massive data warehouses that have lots of  
15 databases or tables, of which there are  
16 applications, that basically they build an  
17 application and then they access the table, right,  
18 and they provision that access.

19 The question is, to get access to that,  
20 right, Facebook probably, the security team, has  
21 some way of tracking who's coming inbound to  
22 Facebook to get data queries for those particular  
23 systems that have user data. At least that's what  
24 I'm thinking. I could be completely wrong here, but  
25 I'm trying to come up with a simple way that --

1 because Facebook already has to secure and track and  
2 monitor to make sure there's not unauthorized access  
3 to user data. So there has to be some sort of form  
4 or function of logging the record keeping that's  
5 kept somewhere about what data is going inbound or  
6 outbound, but only for the narrow subset of systems  
7 as the judge has narrowed and clarified.

8 MS. CLARK: Mr. Garrie, you know, we want  
9 to do everything we can to help you with all this  
10 list, obviously; and if there are specific sources  
11 that you have questions about, we will do our best  
12 to get back to you on those questions as quickly as  
13 we can.

14 What I understand Mr. Pope to be saying is,  
15 you know, how quickly a question could be answered  
16 would depend on the question itself. Some are a bit  
17 more complicated. Some are easier. Some of them  
18 his point of contact might be able to answer. Some  
19 of them his client contact might need to go to  
20 various other people to answer.

21 So I think we'll just need to understand  
22 what the information is you're looking for and for  
23 which source, and from there we can gauge what work  
24 would go into answering the questions.

25 MR. GARRIE: Yeah. And what I'm trying to

1 do is limit the work.

2 So the way I envision this is I'll identify  
3 a subset of systems that he already sent  
4 questionnaires for and got answers to. According to  
5 him, he'd be able to tell me what data could be --  
6 what user data is deleted, does a system have user  
7 data, true or false, and what user data, you know,  
8 needs to be deleted from that system that he said  
9 was in the questionnaire.

10 I haven't seen the questionnaire. So my  
11 goal is to narrow it to a subset of the 149, get the  
12 questionnaire and the answers after counsel's  
13 reviewed them to make sure, then use that to better  
14 understand what we're talking about. Like the  
15 systems he had no idea about, hopefully the  
16 questionnaire will provide sufficient amount of  
17 detail.

18 And then from there, if necessary, I'll  
19 follow up with additional questions again using the  
20 way the judge has defined what is in scope and what  
21 is out of scope.

22 MS. CLARK: I just want to --

23 MS. MUMM: Can I clarify?

24 Yeah. Certainly we will follow up with  
25 whatever questions you have and do our best to

1 answer them, but it is not our understanding that  
2 the questionnaire defines the type of user data.  
3 It's just a yes-or-no question as to whether it may  
4 contain, it can contain information within the scope  
5 of the definition of user data.

6 So the questionnaire will determine whether  
7 it may store user data, but not whether it does and  
8 what type it is. And then the next level is --

9 MR. GARRIE: Well, I think he said that  
10 they also identified. I thought that was the  
11 true-or-false part, and then they also then  
12 identified.

13 That's why I want to see the questionnaire.  
14 I couldn't -- I wasn't exactly sure. I'm trying to  
15 make this as seamless as possible based on what  
16 you've already done because it sounds like to me he  
17 got an answer, and then he went back to them and  
18 said, "What user data is actually going into the  
19 system," and they provided him some information  
20 about it at some level.

21 MS. MUMM: Our understanding is that the  
22 information he has does not relate to what user data  
23 in the sense that you and I would talk about it,  
24 but, rather, how that user data can be, you know,  
25 essentially deleted within the framework; right?

1 The focus is very much on accomplishing that end,  
2 but then -- but again, you know, I think to the  
3 extent you want to narrow the data, we can work with  
4 Mr. Pope and identify what he has. We just don't  
5 want -- we don't want to, you know, set expectations  
6 that are not in line with what we understand the  
7 work task would have done.

8 MS. CLARK: My understanding of the --

9 MR. GARRIE: Yeah.

10 MS. CLARK: -- process.

11 MR. GARRIE: Once I see how he defines what  
12 user data and what answer he got. All I'm trying to  
13 do is avoid having to talk to a hundred -- like,  
14 I'm already convinced I don't have to talk to more  
15 than like 15 or 18 people, but what I'd like to do  
16 is talk to two or three or follow up with just a  
17 handful of people rather than have an extensive  
18 back-and-forth attempt at spreadsheets again.

19 I'd rather just see what's been done  
20 because what I'm not trying to do is make Facebook  
21 go and do a bunch more work to just answer my  
22 question about what -- to answer the judge's  
23 question, frankly, which is pretty salient, which is  
24 what -- and I just need to think about -- I got to  
25 review the transcript, see again what the judge

1 ordered, and then ultimately narrow it down to a  
2 subset of systems for them either to request  
3 Facebook to provide that information about it  
4 because the judge is pretty -- the order is now much  
5 clearer from what it was; right? There's no --  
6 she's removed a lot of the ambiguity there, and then  
7 get that information, and that's how I think I'd  
8 like to proceed.

9 If I can get a rough of the transcript. I  
10 don't need a final rush. A rough is more than  
11 enough just to look at it, to go through it. And  
12 then I will likely issue -- we have a status  
13 hearing. So maybe I'll do it there or I'll issue a  
14 supplemental order. Hopefully we can just do it at  
15 the status hearing accordingly to move it forward  
16 to just get that done based on the judge's  
17 direction.

18 I do like the idea of a status hearing. I  
19 do like that.

20 MS. CLARK: Thank you. And we appreciate  
21 your scheduling that.

22 MR. GARRIE: Counsel Weaver?

23 MS. WEAVER: Thank you. I apologize.

24 We, plaintiffs, would like the opportunity  
25 to review the transcript ourselves and have our

1 experts review it before you issue a final ruling;  
2 and if you could build in a procedure for that, we'd  
3 be grateful.

4 We haven't had an opportunity to weigh in  
5 yet on any of the declarations that have been  
6 submitted or any assertions that are being made by  
7 Mr. Pope or there are other folks who put in  
8 declarations as well.

9 So we would request an opportunity to, at  
10 least, weigh in on these data sources and what we  
11 think they reveal before you issue a final ruling.

12 MR. GARRIE: Okay. So noted for the  
13 record.

14 MS. MUMM: And just to clarify, I don't  
15 have the protective order in front of me, but I  
16 think that there's some limitation on the ability of  
17 experts who are being disclosed at this time, but I  
18 don't have it in front of me.

19 MR. GARRIE: Wait. Wait. Time-out.  
20 Time-out.

21 I don't want -- I assume everybody's  
22 experts are following all the rules that the parties  
23 have set forth and required and defined as; and if  
24 there is any indication that hasn't been the case,  
25 feel free to raise the issue separately.



1 I will take your issue, Counsel Weaver,  
2 under advisement and think about it. I just want to  
3 read the rough transcript, read her order again,  
4 read the transcript again, and then come up with a  
5 game plan and roadmap that's actually feasible  
6 because she issued three or four -- four orders.  
7 So I'm trying to do all of that and move it all  
8 along.

9 So that's kind of where my headspace is.  
10 I'm not opposed to hearing from your experts or  
11 having them weigh in. I just want to make it in a  
12 methodical way so you don't have to do it more than  
13 once and Facebook doesn't have to respond more than  
14 once. I'd rather get forward until that point.

15 And then rather than issuing a formal  
16 order, can you ask Mr. Pope for an empty  
17 questionnaire or blind questionnaire that he sent  
18 out and just circulate it?

19 MS. CLARK: My understanding, and Ms. Mumm  
20 can chime in a little more here, is that there is  
21 not a single questionnaire, that the questionnaires  
22 are tailored to a team.

23 MS. MUMM: Right. That's what he said  
24 during his questioning and that was our  
25 understanding.

1           So if you can specify the systems you're  
2     interested in, we can get a specific questionnaire  
3     for each.

4           MR. GARRIE:   Okay.   And what about just  
5     what user -- I mean, he had to use the same  
6     definition of user data.

7           MS. MUMM:   We can follow up with that.   I  
8     think that that might be --

9           MS. WEAVER:   One --

10          MR. GARRIE:   Because if it turns out --

11          MS. WEAVER:   Go ahead.

12          MR. GARRIE:   Go ahead, Counsel Weaver.

13          MS. WEAVER:   One question that we on the  
14     plaintiffs' side had is the timing for these data  
15     sources.   It sounded as though the testimony that  
16     came in today, that this list was prepared maybe a  
17     year ago and we don't know its scope.   Certainly  
18     our class period began in 2007.   I don't know if  
19     this is what was prepared and if it was a year ago,  
20     whenever it was, but it is a question to be  
21     answered.

22                 Are there other data sources that are not  
23     active that might contain older data for some of the  
24     plaintiffs?   I just don't know because this was not  
25     collected for this case.   It was collected for

1 another purpose.

2 MR. GARRIE: No. I agree. It's noted for  
3 the record.

4 Counsel Mumm, did you want to respond?

5 MS. MUMM: Yeah. I'll just note for the  
6 record that our ESI protocol indicates that we're  
7 not required to collect or produce some off-line  
8 statements questions.

9 MR. GARRIE: Noted for the record and not  
10 lost on the special master.

11 The other question I had is I got your  
12 post, Counsel Kutscher Clark, and I just want to  
13 make sure. The Strow's (phonetic) FTI report,  
14 you're going to start producing the report; right?  
15 I just -- your counsel, Alex. I forget his last  
16 name. Counsel Alex is in the process of reviewing a  
17 report for privilege, but there will be a rolling  
18 production starting when exactly? I didn't see a  
19 date.

20 MS. CLARK: I'm sorry. I misunderstood.  
21 You were referring to Alex Cadwell.

22 We're working on it currently and we expect  
23 we will start the production next week. We're in  
24 the process to get them collected and into the  
25 database, but that is all underway and the wheels

1 are in motion and they should start going out the  
2 door next week.

3 MR. GARRIE: Okay. And I have no problem  
4 with the reconsideration. That's fine. I will  
5 respond to plaintiffs' request to be heard on their  
6 motion for reconsideration.

7 Any other issues plaintiffs or Facebook  
8 want to raise? And I declare an impasse. Well,  
9 spoke to Judge Andler and we moved that along as  
10 well. I posted that, I believe, yesterday or this  
11 morning or I emailed at some point the impasse on  
12 those two issues.

13 So, I guess, Plaintiffs, anything you'd  
14 like to cover since I have everybody here?

15 MS. WEAVER: I think that's it. Just to be  
16 clear, it's not so much responding to the motion for  
17 reconsideration as to have an opportunity to weigh  
18 in on the collection of data sources and in light of  
19 Judge Corley's order what's within the scope of  
20 discovery.

21 MR. GARRIE: Yeah. I fully -- yeah. I  
22 understand.

23 MS. WEAVER: All right.

24 MR. GARRIE: Counsel for Facebook?

25 MS. CLARK: I think we are all set. Thank

1     you.

2                   MR. GARRIE: All right. We'll go off the  
3     record.

4                                   \* \* \*

5           (Whereupon, the hearing ended at 3:21 p.m. PST)

6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25

CERTIFICATION OF COURT REPORTER

FEDERAL JURAT

I, the undersigned, a Certified Shorthand Reporter of the State of California do hereby certify:

That the foregoing proceedings were taken before me at the time and place herein set forth; that any witnesses in the foregoing proceedings, prior to testifying, were placed under oath; that a verbatim record of the proceedings was made by me using machine shorthand which was thereafter transcribed under my direction; further, that the foregoing is an accurate transcription thereof. I further certify that I am neither financially interested in the action nor a relative or employee of any attorney of any of the parties.

IN WITNESS WHEREOF, I have this date subscribed my name: Date: January 25, 2022.



Michelle Milan Fulmer

CSR 6942, RPR, CRR, CRC

[&amp; - additional]

<b>&amp;</b>	<b>29</b> 122:14	<b>absolutely</b> 71:12	97:5 99:8 100:8
<b>&amp;</b> 3:4,17,21 36:18 36:22 46:25 47:6	<b>3</b>	<b>accelerate</b> 68:6	100:23 101:15
<b>0</b>	<b>3</b> 7:12 13:25	<b>acceptable</b> 73:23	102:7 106:14
<b>0</b> 42:10	<b>3200</b> 3:12	<b>access</b> 18:3 21:22	107:18 111:18
<b>02843</b> 1:6 2:6	<b>351-2404</b> 3:23	21:24 24:3,4,8,16	112:13 113:2,4
<b>1</b>	<b>3:21</b> 6:2 135:5	26:5 42:23 43:12	114:20 122:20,21
<b>1</b> 4:3	<b>4</b>	43:14,15,22 45:9	<b>accessing</b> 28:1,15
<b>10</b> 40:15,17 123:16	<b>415</b> 3:7	47:18 52:9 53:6	28:16 70:7,8
<b>10166-0193</b> 3:23	<b>445-4003</b> 3:7	53:16 54:4,7,10,11	<b>accomplishing</b>
<b>10th</b> 6:23	<b>5</b>	55:4,9,22 58:2,4,7	128:1
<b>1201</b> 3:12	<b>555</b> 3:6	58:9 59:23 60:4	<b>account</b> 56:9 81:6
<b>12:08</b> 2:12 6:2	<b>6</b>	64:21 72:8 85:7	81:7
<b>12th</b> 3:6	<b>6</b> 5:5	91:24 92:14 93:15	<b>accurate</b> 23:3 30:6
<b>14</b> 1:15 2:13 6:1	<b>623-1900</b> 3:13	94:4 95:8,15,19	40:2 48:8 56:22
<b>149</b> 9:15 18:17	<b>650</b> 3:19	98:21 100:9	66:2,3 77:5 101:7
123:21 124:10	<b>6942</b> 1:18 136:25	106:25 107:2,23	107:4 136:14
126:11	<b>7</b>	108:8 110:3,12	<b>acquired</b> 36:24
<b>15</b> 123:16 128:15	<b>71</b> 5:11	112:17 113:15,21	47:6 91:1
<b>15030</b> 136:23	<b>8</b>	113:24,25 114:4	<b>acquisition</b> 35:16
<b>1600</b> 3:6	<b>849-5348</b> 3:19	116:5,16 119:21	<b>acquisitions</b> 36:23
<b>18</b> 1:6 2:6 128:15	<b>9</b>	124:17,18,19	<b>acronym</b> 34:21
<b>1881</b> 3:18	<b>94025</b> 4:3	125:2	45:22
<b>1:26</b> 66:12	<b>94304-1211</b> 3:18	<b>accessibility</b> 49:5	<b>action</b> 62:25
<b>1:29</b> 68:13	<b>94607</b> 3:6	<b>accessible</b> 12:21	115:22 136:16
<b>1:46</b> 68:14	<b>98101</b> 3:12	13:8 16:9 21:21	<b>actions</b> 63:11
<b>2</b>	<b>a</b>	26:6,8 27:12	<b>active</b> 77:20,22,23
<b>20</b> 108:23	<b>ability</b> 53:6 113:4	30:21 32:13 35:13	83:17 132:23
<b>200</b> 3:22	124:4 130:16	35:22 37:18 39:13	<b>activity</b> 25:13
<b>2007</b> 90:2 132:18	<b>able</b> 11:10 12:14	39:25 41:4 42:6	<b>actor</b> 62:25
<b>2014</b> 75:25	33:24 56:9 64:16	46:6 50:2 51:9	<b>actual</b> 24:13 32:8
<b>2021</b> 6:23	80:10 87:24 88:21	56:24 57:6,7,21	79:2 80:13 81:25
<b>2022</b> 1:15 2:13 6:1	89:6 97:11 98:16	60:2,23 61:9,14,15	87:19 88:24
136:19	113:14 114:15	62:7 63:3 64:1,8	<b>ad</b> 13:7,11,12
<b>206</b> 3:13	117:15 125:18	64:22 65:2,20	14:14 16:18 74:10
<b>212</b> 3:23	126:5	67:5,12,13,24	123:25
<b>24</b> 52:18	<b>abp</b> 10:21 11:11	69:17 70:20 72:24	<b>adavis</b> 3:8
<b>25</b> 136:19	12:8	72:25 73:13 75:10	<b>additional</b> 15:3
<b>2843</b> 1:5 2:5		75:20 77:13,14	67:17,20 94:11
		81:15 85:4,20	117:19 121:15
		87:4 89:9,16 90:5	123:14 126:19
		90:16 96:13,20	

## [address - assessments]

<b>address</b> 71:14 <b>addresses</b> 71:18 <b>addressing</b> 33:11 <b>adfox</b> 14:8,22 15:14 <b>adlogger</b> 16:7,10 <b>admarket</b> 17:8 <b>admarketdb</b> 27:9 <b>admission</b> 89:1 <b>adoption</b> 110:25 <b>ads</b> 18:8,8 26:3 27:10,14,17 30:20 <b>advertisers</b> 26:2 74:19 <b>advise</b> 84:14 <b>advisement</b> 131:2 <b>afternoon</b> 6:15 <b>age</b> 23:4 <b>aggregate</b> 16:19 35:14 <b>aggregated</b> 18:10 80:21 89:18 111:23 <b>aggregates</b> 111:25 <b>aggregating</b> 90:1 <b>aggregation</b> 89:25 <b>ago</b> 77:8,24 111:1 132:17,19 <b>agree</b> 59:16 98:13 133:2 <b>ahead</b> 9:11 11:17 11:18 20:12 38:18 51:19 71:17 79:16 132:11,12 <b>ai</b> 55:22 115:3 <b>aka</b> 38:9 <b>akkio</b> 27:21,25 <b>alex</b> 133:15,16,21 <b>allocated</b> 62:18 <b>allow</b> 110:11	<b>allows</b> 32:23 97:18 <b>alphabetical</b> 114:16 <b>alto</b> 3:18 <b>amazon</b> 67:9 <b>ambiguity</b> 129:6 <b>ambiguous</b> 89:25 <b>amount</b> 109:17,17 126:16 <b>analogous</b> 67:9 <b>analysis</b> 83:15 89:15 94:15 95:4 99:1 <b>analytics</b> 35:11 46:17 <b>analyzing</b> 77:24 94:22 <b>andler</b> 134:9 <b>anne</b> 3:5 <b>anonymi</b> 82:19 <b>anonymized</b> 44:23 78:4,6,6,16 79:2,3 79:4 80:22 81:16 <b>answer</b> 6:17 11:3 11:7 17:15 21:9 27:17 45:11 55:23 68:2 72:2 75:7 82:24 83:7 84:16 84:19 87:6 88:14 98:16 119:14 123:25 124:5 125:18,20 127:1 127:17 128:12,21 128:22 <b>answered</b> 125:15 132:21 <b>answering</b> 123:4 125:24 <b>answers</b> 17:24 123:20 126:4,12	<b>anybody</b> 33:1,13 56:24 93:11,15 107:18 117:24 <b>anyways</b> 38:20 75:8 <b>apart</b> 122:23 <b>api</b> 85:7 113:19,24 114:4 <b>apis</b> 113:25 <b>apologize</b> 66:11 129:23 <b>app</b> 28:10 31:22 34:13,16,21 37:8 37:14 51:24 53:4 53:5 54:9 58:7,8 59:11 60:3 72:22 74:20 75:10 85:8 89:23 95:14 105:11 107:2,17 110:4,11 112:14 <b>appearance</b> 71:19 <b>appearances</b> 3:1 <b>apples</b> 80:11,11 84:23,23 <b>application</b> 23:23 24:1,3 26:18 35:2 37:14 44:7 52:7,7 53:12,17 54:4 58:11 65:14 86:24 87:23 89:19 97:13 100:1 106:6 107:2 108:13,15 109:1 109:14,16 110:8 119:22 124:17 <b>application's</b> 65:11 78:15 <b>applications</b> 24:4 24:7,18 33:8 37:25 44:3 50:2 56:13 58:3 65:10 65:15,15 68:6,7,8	82:5 87:17 88:24 94:1 102:22 108:11 109:11 112:8 115:13 119:20 124:16 <b>appreciate</b> 17:1,14 117:14 120:21 121:20 129:20 <b>appreciative</b> 118:13 <b>appropriate</b> 15:3 29:25 73:24 118:9 119:24,24 123:2 <b>apps</b> 79:9 <b>ar</b> 37:2 61:24 62:2 <b>archival</b> 30:25 98:8 <b>area</b> 14:19 22:6 101:9 102:25 <b>artifacts</b> 10:10 <b>artificial</b> 55:18 <b>artillery</b> 31:12,13 <b>aside</b> 64:19 70:14 76:21 <b>asked</b> 8:15 9:9,12 10:7 37:11 78:12 83:14,23 86:21 92:5 119:16 120:11 <b>asking</b> 8:9 15:16 41:11 55:19 62:6 66:7 67:4 71:5 80:17 84:15,22 101:3 104:22 109:24 123:8 <b>asks</b> 89:1 <b>assert</b> 110:1 <b>assertions</b> 130:6 <b>assessments</b> 48:6 48:19
---	---	--	---



[asset - bytes]

<b>asset</b> 39:9,22 <b>assistant</b> 32:12 <b>associated</b> 82:16 <b>assume</b> 13:18 16:12 23:15 25:24 27:17 32:17 38:3 39:24 42:12 54:14 60:22 63:23 65:1 69:23 71:24 78:20 85:19 88:10 92:13 95:9,18 96:19 97:19 98:14 100:12 101:6,19 114:2 116:12 130:21 <b>assuming</b> 36:19 79:1,22 81:2 96:14 119:19 <b>assumption</b> 38:5 40:2 48:8 64:3 77:2 101:7 107:4 114:6 <b>async</b> 32:17,21,22 33:23 41:6 57:3 <b>attached</b> 7:19 <b>attempt</b> 128:18 <b>attorney</b> 136:17 <b>attorneys</b> 1:10 120:14 <b>audience</b> 26:3 42:5 <b>augmented</b> 61:24 <b>auld</b> 3:4 <b>availability</b> 102:14 <b>available</b> 10:2 17:23 20:16 21:2 22:3 28:12,25 30:21 31:18,22 32:13 33:1 35:25 40:7 44:9 45:24	46:6,21 48:7 50:15 56:1 61:9 62:7 65:2,19,20 72:13,22,23 73:13 74:19 75:9 78:3 81:14 85:4,19,20 87:3,4 89:14 90:5 90:15 91:12 93:11 96:13 100:23 105:11 106:12,13 111:18 114:13 118:19,24 119:2,3 119:5,8,9 <b>avenue</b> 3:12,22 <b>aviv</b> 20:20 21:21 22:12 <b>avoid</b> 128:13 <b>aware</b> 6:20 22:5 25:23 <b>aws</b> 10:21 11:20 11:21,22,22,23 58:15	<b>backup</b> 31:2,11 42:11 47:16 <b>backups</b> 9:25 95:9 98:9 <b>balance</b> 33:25 <b>bars</b> 34:12,13,16 <b>based</b> 11:1 23:7 33:8 50:22 88:18 102:1 127:15 129:16 <b>basic</b> 11:2 <b>basically</b> 8:15 9:23 35:1 36:14 43:8 49:1 50:12 52:24 55:4 68:22 100:1 110:16 124:16 <b>batch</b> 96:25 <b>batched</b> 33:22 <b>beat</b> 35:10 <b>beest</b> 36:3 <b>began</b> 132:18 <b>beginning</b> 60:20 <b>behavior</b> 50:25 <b>believe</b> 35:16 59:4 69:20 84:11 134:10 <b>bento</b> 36:7 <b>best</b> 40:13 54:12 76:17 125:11 126:25 <b>better</b> 126:13 <b>beyond</b> 24:17 32:11 50:7 63:22 64:10 75:1 76:9 76:24 <b>bfalaw.com</b> 3:7,8 3:8 <b>bi</b> 98:24 <b>big</b> 15:2 40:15,16 <b>bigbox</b> 36:18	<b>binaries</b> 116:21 <b>binary</b> 8:23 32:10 56:17 103:17 119:16 <b>bit</b> 19:4 23:17 24:22 27:22 32:20 43:6 52:5 56:6 83:8 103:6 125:16 <b>bleichmar</b> 3:4 <b>blind</b> 131:17 <b>blob</b> 19:19 20:14 49:21 <b>block</b> 55:4 <b>bottom</b> 55:12 81:10 <b>box</b> 109:6 <b>break</b> 66:14 <b>breaks</b> 46:3 <b>breathing</b> 81:8 <b>bring</b> 117:7 121:5 121:6 <b>broad</b> 80:1 83:8 108:7 <b>broader</b> 29:1 <b>broke</b> 10:19 <b>build</b> 23:6 42:5 54:6,9 113:7 124:16 130:2 <b>building</b> 70:8 89:22 <b>builds</b> 53:4 <b>built</b> 38:24 39:2 42:16 45:19 55:24 75:15 95:14 <b>bunch</b> 94:25 95:1 128:21 <b>busy</b> 6:17 <b>bytes</b> 72:17
	<b>b</b> <b>b</b> 5:8 <b>back</b> 11:13 27:7 39:20 41:23 42:4 42:7,7 43:23 51:15 68:13 71:22 80:8,23 82:20,24 84:17,18 86:19 90:2 96:6 105:24 108:18 117:7 118:1 121:5,7 123:19 125:12 127:17 128:18 <b>backbone</b> 27:10 <b>background</b> 33:23 <b>backing</b> 41:6 102:9,10 <b>backs</b> 23:21		

[c - comparable]

<b>c</b>	132:25	<b>choice</b> 72:18	<b>clock</b> 66:10
<b>c</b> 3:22	<b>cases</b> 29:8,10,11	<b>chronos</b> 24:20,21	<b>close</b> 20:8 83:25
<b>cache</b> 63:11 64:25	29:25 30:4,8,11,14	24:24 25:8,11,17	<b>closer</b> 19:23 28:18
76:19,22,25 81:12	30:19 60:14 78:14	32:18 37:11	<b>clue</b> 11:20
81:13 102:22	96:3 99:22 106:18	<b>circulate</b> 131:18	<b>coast</b> 28:18,19
<b>caches</b> 19:16,20	<b>cat</b> 20:5,13 21:10	<b>clarification</b>	66:12
40:6	<b>category</b> 69:19	123:19 124:11	<b>cocoon</b> 25:22 26:6
<b>caching</b> 102:11	<b>cause</b> 83:15	<b>clarified</b> 125:7	37:12
<b>cadwell</b> 133:21	<b>ccpa</b> 15:25	<b>clarify</b> 33:15,16	<b>code</b> 33:4 54:14
<b>calculating</b> 106:9	<b>cdn</b> 18:19,22 19:5	38:23 57:25 78:24	55:5,5 59:4 94:2
<b>california</b> 1:2 2:2	19:6,12,12,16 20:7	81:22 87:13 90:23	114:2
2:12 3:6,18 4:3	20:14,19 21:2,8,25	93:25 97:11 105:8	<b>coding</b> 103:25
136:5	22:2,7,20 37:11	115:7 126:23	<b>coefficient</b> 26:11
<b>call</b> 18:19 46:2	40:5	130:14	26:12,13 37:12
89:22,23,25 90:1	<b>center</b> 28:18	<b>clarity</b> 84:19	<b>collab</b> 37:13,13,14
94:22	<b>central</b> 44:4	123:14	<b>collaboratively</b>
<b>called</b> 6:9 20:22,23	115:15	<b>clark</b> 3:17 71:4,12	37:15
37:24 45:19 80:19	<b>centrally</b> 115:18	71:14,20 80:12,15	<b>collect</b> 8:2,7 79:6
112:3	<b>certain</b> 28:10,10	84:9 120:10 121:9	80:2 118:18 133:7
<b>callisto</b> 23:9 24:18	<b>certainly</b> 17:25	125:8 126:22	<b>collected</b> 12:19,20
37:10 51:16	74:7 118:10	128:8,10 129:20	32:1 123:7 132:25
<b>calls</b> 89:24	126:24 132:17	131:19 133:12,20	132:25 133:24
<b>calypso</b> 22:22	<b>certification</b> 136:1	134:25	<b>collecting</b> 44:24
37:11	<b>certified</b> 2:11 6:10	<b>class</b> 132:18	<b>collection</b> 134:18
<b>campen</b> 3:11	136:4	<b>classification</b>	<b>collects</b> 123:3,5
<b>candidate</b> 97:9	<b>certify</b> 136:6,15	30:16 32:5	<b>combined</b> 72:1,3
<b>capabilities</b> 88:4	<b>challenge</b> 120:18	<b>classified</b> 8:22	<b>combines</b> 10:25
92:19	<b>chance</b> 18:16	<b>classify</b> 9:5	<b>come</b> 30:8 93:7
<b>capability</b> 29:5	38:13 44:25	<b>claufenberg</b> 3:14	124:25 131:4
<b>capitalize</b> 85:24	101:14 120:4	<b>clean</b> 10:24	<b>comes</b> 22:25
<b>caps</b> 15:2	<b>channel</b> 22:4	<b>cleanup</b> 25:10,12	<b>coming</b> 124:21
<b>captured</b> 65:7	<b>channels</b> 84:5	<b>clear</b> 12:20 15:5	<b>commencing</b> 2:12
<b>captures</b> 96:11	<b>chat</b> 61:21 71:15	16:7 24:10 35:12	<b>comment</b> 20:20
<b>capturing</b> 30:16	71:18	48:11 51:23 67:22	<b>common</b> 69:5
<b>care</b> 12:6	<b>check</b> 8:23 17:20	74:18 101:12	103:23
<b>cari</b> 3:11	17:23 89:2,2	124:3 134:16	<b>company</b> 18:1
<b>case</b> 1:6 2:6 6:20	109:6	<b>clearer</b> 129:5	36:24 76:16 91:1
20:3 29:4,4 30:9	<b>chen</b> 4:2 84:13	<b>click</b> 109:6	95:20 98:7 115:10
45:20 50:18 54:19	117:24	<b>clicks</b> 74:11	115:20 116:24,25
69:5,7 93:6 109:7	<b>chime</b> 131:20	<b>client</b> 84:11	<b>comparable</b> 69:3
109:25 130:24		125:19	

**[competent - cycle]**

<b>competent</b> 84:13 <b>compile</b> 94:2 <b>completely</b> 25:5,7 36:5 124:24 <b>compliance</b> 48:6 48:19 115:3 <b>compliant</b> 49:18 49:19 <b>complicated</b> 125:17 <b>component</b> 41:6 107:13 <b>composability</b> 98:17 <b>compression</b> 66:5 <b>compute</b> 46:16 <b>computer</b> 7:6 <b>concedes</b> 122:25 <b>concern</b> 124:2 <b>concerned</b> 105:4 <b>confer</b> 84:18 <b>confidential</b> 1:9 40:22 120:14,15 <b>confidently</b> 65:5 <b>configerator</b> 37:20 37:21,23,24 <b>configuration</b> 28:11 32:6 37:24 88:4 89:3 <b>configure</b> 9:5 29:4 100:1 109:19 <b>configuring</b> 15:19 30:17 <b>confirm</b> 32:19 41:3 42:13 46:10 48:9 60:11,24 64:3,19 65:3,5 70:11 74:16 77:16 78:22 96:20 98:4 <b>confusing</b> 16:15	<b>connected</b> 39:5 <b>connection</b> 6:23 <b>connectivity</b> 38:2 <b>consider</b> 59:22 82:14 <b>consumer</b> 1:6 2:6 53:2 76:10 <b>consumers</b> 18:7 <b>contact</b> 14:25 15:15 17:5 87:8 87:15,20 88:1 125:18,19 <b>contacting</b> 87:22 <b>contacts</b> 73:18 <b>contain</b> 76:23 127:4,4 132:23 <b>contained</b> 11:4 <b>container</b> 102:1 <b>containers</b> 102:2 <b>contemporaneous</b> 52:14 <b>content</b> 18:23 19:14 22:17 37:15 63:18 82:10 <b>context</b> 75:3 98:19 105:3,21 <b>contextual</b> 86:9 <b>contributor</b> 38:9 38:11 39:1 <b>contributors</b> 38:10 <b>control</b> 68:23 69:4 76:8 89:1 <b>controls</b> 113:25 <b>conversation</b> 120:4 121:2 123:18 <b>conversations</b> 44:5 <b>conversions</b> 74:13	<b>convey</b> 104:11 <b>convinced</b> 128:14 <b>copied</b> 20:14 <b>copies</b> 19:17 <b>copy</b> 19:23 78:23 <b>core</b> 16:14,15 27:10 41:2 116:23 116:23 <b>corley's</b> 134:19 <b>correct</b> 7:23 17:10 21:4 25:15 33:3 33:16 37:23 38:4 40:14 42:18,22 47:20 48:4,22 49:20,20 54:8,13 55:10,14 56:2,23 57:1 60:21 61:22 67:10 68:8 70:22 75:19 79:5 97:4 107:6,19,19 <b>correctly</b> 41:5 <b>cortext</b> 39:9,21 <b>cosco</b> 39:24 <b>counsel</b> 3:1 7:1 71:17 80:11 81:19 84:10,12,13,13,18 117:7,9,24 118:1,3 120:6 121:6,8 122:11 129:22 131:1 132:12 133:4,12,15,16 134:24 <b>counsel's</b> 126:12 <b>counters</b> 100:2 <b>couple</b> 27:8 39:6 <b>coupled</b> 52:8 <b>course</b> 8:1 35:9 <b>court</b> 1:1 2:1 6:6 39:16,18 40:22,23 70:25 71:16 120:17 122:23	123:1,2 136:1 <b>court's</b> 120:5 121:4 <b>cover</b> 121:19 123:21 134:14 <b>covered</b> 37:10,11 69:3 83:7 <b>covers</b> 21:7 69:20 <b>crazy</b> 17:13 <b>crc</b> 1:18 136:25 <b>create</b> 98:1 111:5 <b>created</b> 23:10,11 23:12 99:17 <b>creating</b> 37:15 <b>credential</b> 81:7 <b>criteria</b> 80:4 <b>critical</b> 18:7 <b>cross</b> 81:3 <b>crr</b> 1:18 136:25 <b>crutcher</b> 3:17,21 <b>crypto</b> 48:2,21 97:2 <b>cryptocurrency</b> 47:24 48:14,16,18 <b>csr</b> 1:18 136:25 <b>cubrick</b> 40:10 <b>cumbersome</b> 87:14 <b>current</b> 10:6 <b>currently</b> 133:22 <b>customer</b> 116:12 116:13 <b>customer's</b> 88:24 <b>customers</b> 14:13 23:6 48:15 <b>cut</b> 95:1 <b>cws</b> 41:2 <b>cycle</b> 14:3
---	---	--	--

[d - deleted]

<b>d</b>	54:6,17,21,24 55:5	109:13,16,20	<b>db</b> 91:20 96:23
<b>d</b> 5:1 42:4	55:10 56:19 57:6	110:9,12 111:19	<b>de</b> 78:6
<b>daily</b> 77:20,22,23	57:14,22,24 58:5,7	111:21 112:4,12	<b>deal</b> 78:13
<b>daiquery</b> 41:18	58:18 59:21 60:4	112:25 113:5,17	<b>december</b> 6:23
<b>dam</b> 39:10	60:12,17,23 61:7	114:3,20,22 116:3	<b>decision</b> 29:23
<b>daniel</b> 1:11 54:6	61:15 63:2,6,15,16	116:4 118:16,18	<b>declaration</b> 6:19
<b>daniel's</b> 58:8	63:25 64:4,14	118:20 119:2,3,6	6:22 7:3,5,8,11,12
<b>dap</b> 77:7,20	65:6,9,13 66:6	119:17,17,19	17:20 26:22 71:3
<b>dapl</b> 42:4	67:11,14,18,23	121:17,19 122:18	71:5 81:1 86:21
<b>dark</b> 42:10	68:16,24 69:1,7	122:24,25 123:3,5	<b>declarations</b> 130:5
<b>data</b> 7:13,14 8:8	70:7,9,16,18 72:16	124:1,14,22,23	130:8
8:17,18,22,22,24	72:19,24 73:10,12	125:3,5 126:5,6,7	<b>declare</b> 134:8
8:24 9:2,3,5,7,25	74:9,10,24 75:6,9	126:7 127:2,5,7,18	<b>dedicated</b> 19:13
10:2,13,23,23 11:3	75:12,21,22 76:18	127:22,24 128:3	<b>deemed</b> 29:24
11:12,25 12:1,2,4	76:19,23 77:13	128:12 130:10	<b>deep</b> 94:22
12:7,18,19,20,20	78:5,5,6,7,8,12,16	132:6,14,22,23	<b>defendant</b> 3:16
13:4,8,15,16 14:2	78:20,25 79:1,2,2	134:18	4:1
14:4 15:5,7,10,16	79:3,4,5,10,18,25	<b>database</b> 9:1,2	<b>define</b> 38:11 83:11
15:18,23,24 16:8	80:13,19,19,20,21	16:12 17:8,9	<b>defined</b> 78:25
16:16 18:9 19:15	80:21,22 81:2,4,11	25:10,11,12 54:4,7	112:14 126:20
19:17 21:9,15,19	81:11,13,21,25	69:21 102:10	130:23
21:25 22:2,25	82:6,8,9,11,11,12	107:2 108:9,12	<b>defines</b> 127:2
24:20 25:1,1 26:4	82:12,14,18 83:12	109:3 110:12,15	128:11
26:7,7,13,15,19,22	83:16 86:15,17,18	110:17,19 112:5	<b>definitely</b> 102:25
26:25,25 27:1	86:24,25 87:16,18	113:2,6 133:25	111:22
28:1,10,13,17,17	87:19,21 88:9,11	<b>databases</b> 41:23	<b>definition</b> 79:25
28:17,20,22,24,25	88:18,22,25 89:1,4	110:18 124:15	81:24 127:5 132:6
28:25 29:1 30:16	89:8,13,16,16	<b>dataset</b> 54:11	<b>definitive</b> 103:6
31:25 32:5,7,8,11	90:22 91:3,3,11,12	113:21	<b>definitively</b> 60:11
32:14 33:7,10,12	91:17 92:5 93:1	<b>dataswarm</b> 42:16	65:18,24 77:2
35:21 36:15 37:16	94:2,15,23,23 95:1	42:18	<b>degree</b> 25:18
37:16,17 38:20	95:3,25 96:1,15,15	<b>date</b> 133:19	<b>delayed</b> 32:23
39:2,12,25 40:1,6	96:23,25 98:9,23	136:18,19	<b>deld</b> 43:1,1,2,3,4
40:7 41:3 43:9	98:25 99:9 100:9	<b>dated</b> 103:6	107:9,13
44:6,10,10,18,21	100:23,24 102:19	<b>david</b> 5:3 6:8	<b>delete</b> 44:5 79:4
44:22,23,24 45:9	103:17,17 104:2,6	<b>davis</b> 3:5	89:3 108:23 109:8
45:10,16 46:16	104:6,7,15,16,23	<b>dawn</b> 90:19,25	109:12,20,21
47:19 48:10,16	104:24 105:5	<b>day</b> 6:17	110:2
49:12,24,25 50:8	106:12,22 107:1,1	<b>days</b> 77:8,24	<b>deleted</b> 21:14 44:6
50:10,13,15 52:16	107:3 108:16,18	109:18,20,22	52:12 92:7 126:6
52:17,19 53:20	108:23,24 109:6,9		126:8 127:25

[deletes - either]

<b>deletes</b> 43:9,10,17 <b>deletion</b> 9:6,8 10:1 13:25 14:1 15:20 30:17 32:6 43:8 43:18,24 44:4 53:10 88:4 107:8 107:11 109:9 <b>deletions</b> 107:11 <b>delivered</b> 19:25 <b>delivering</b> 56:12 <b>delivery</b> 14:9,11 18:8,23 19:13 24:14 44:15 45:2 45:3 64:7,14 <b>depend</b> 125:16 <b>dependency</b> 45:7 <b>deployed</b> 110:23 <b>deployment</b> 58:25 <b>deprecated</b> 99:17 99:18 <b>depth</b> 12:15 25:18 29:6 36:2 52:5 <b>derek</b> 3:11 <b>described</b> 22:12 29:17 63:22 64:17 <b>describing</b> 14:17 52:22 55:6 <b>description</b> 5:10 9:24 14:9 16:10 25:2 29:19 31:17 46:4,15 53:23 64:11 65:8 72:6 73:14 98:6,11 101:10 104:10 108:7 114:24 <b>designation</b> 120:22 <b>designations</b> 120:18 <b>designed</b> 41:12 85:1	<b>detail</b> 9:15 87:12 91:10,14 100:11 126:17 <b>details</b> 8:24 16:21 30:10 71:6 74:6 <b>detects</b> 28:15 <b>determination</b> 65:12 <b>determine</b> 122:12 127:6 <b>determined</b> 109:17 <b>determines</b> 28:14 <b>determining</b> 122:13 <b>dev</b> 45:14 98:4 <b>devdb</b> 45:13 <b>developed</b> 108:12 <b>developer</b> 53:4,6 60:3 76:11 85:8 97:13 109:1 110:4 <b>developers</b> 31:22 53:16 55:9 69:10 69:10 70:7 72:22 74:20 75:10 93:11 97:25 100:6 105:11 110:12 <b>developing</b> 47:25 <b>development</b> 31:14 35:6,7 <b>device</b> 56:9 <b>devices</b> 56:8 <b>devops</b> 45:17,20 46:5,7 <b>dictionaries</b> 66:5 <b>differ</b> 100:18 <b>different</b> 16:12 40:11 49:15 55:2 62:19 71:10 82:3 83:19 88:13,18 95:2,3 100:15	107:21 112:8 <b>differentiate</b> 105:14 <b>digital</b> 39:9,22 <b>digraph</b> 46:14 <b>direct</b> 43:12,14 98:10 113:20 <b>direction</b> 129:17 136:13 <b>directly</b> 10:4 12:13 24:18 52:8 60:1 62:3 113:15 117:11 <b>disaster</b> 49:19 <b>disclosed</b> 130:17 <b>discoverable</b> 61:8 65:21 122:18 <b>discovery</b> 122:24 134:20 <b>discrete</b> 82:20 <b>discuss</b> 30:14 <b>discussed</b> 29:13 44:5 <b>discussing</b> 30:15 53:11 <b>discussion</b> 30:19 <b>dispute</b> 6:24 <b>distinction</b> 83:6 100:16 104:8 105:1,3 <b>distributed</b> 19:21 115:19 <b>distributes</b> 20:19 <b>district</b> 1:1,2 2:1,2 <b>dit</b> 45:18,19 <b>diving</b> 94:22 <b>division</b> 91:2 116:24 <b>dloeser</b> 3:13 <b>document</b> 9:4 71:10	<b>doing</b> 17:10 29:16 42:3 87:2 93:6 94:15 96:4 105:4 <b>door</b> 134:2 <b>dot</b> 81:3 85:11,14 <b>double</b> 114:18,21 <b>download</b> 20:22 20:24 <b>downpour</b> 46:24 <b>downstream</b> 89:10,12 <b>dr</b> 49:18 <b>drive</b> 49:2 <b>driving</b> 27:4 102:22 <b>duly</b> 6:9 <b>dumbo</b> 47:13 <b>dunn</b> 3:17,21 <b>dx</b> 47:21 <b>dyi</b> 20:25 21:1,3,6 113:1
<b>e</b>			
<b>e</b> 5:1,8 <b>earlier</b> 69:4 70:3 110:7 <b>early</b> 110:25 122:9 <b>easier</b> 125:17 <b>east</b> 28:18 <b>easy</b> 112:10 <b>ee</b> 48:12,24 98:6 <b>effect</b> 75:24 <b>effectively</b> 15:24 <b>effort</b> 11:22 87:9 87:15 99:17 <b>efforts</b> 112:3 <b>either</b> 62:16 63:10 64:6 67:13 69:14 76:1 85:14,23 86:14 97:16 118:22 129:2			



[elasticsearch - facebook]

<b>elasticsearch</b> 49:3 49:10,11,11 <b>elephant</b> 47:15 <b>email</b> 3:7,13,19,24 4:4 9:16 71:14,18 <b>emailed</b> 92:4 134:11 <b>embedding</b> 115:1 <b>employee</b> 28:4,8 76:13 110:5 115:25 136:16 <b>employees</b> 36:14 41:16,17 72:7 76:17 93:22 94:20 97:25 98:8 115:21 <b>empty</b> 131:16 <b>encoded</b> 103:24 104:18 <b>ended</b> 122:8 135:5 <b>engagement</b> 30:14 <b>engine</b> 46:16 60:6 <b>engineers</b> 55:20 55:21 <b>english</b> 104:17 <b>enjoy</b> 121:13,24 <b>ensuring</b> 13:22 <b>enter</b> 71:1 <b>entire</b> 71:2 79:15 107:12 108:22 <b>entity</b> 13:17 81:5,8 82:21 <b>entrepreneurial</b> 69:25 <b>envision</b> 126:2 <b>esi</b> 133:6 <b>esq</b> 3:4,5,5,11,11 3:17,22 4:2 <b>essentially</b> 18:2 115:9 127:25 <b>estimation</b> 25:24	<b>event</b> 16:13 46:1,4 <b>events</b> 46:10 71:24 <b>everstore</b> 49:9,18 67:16,16,19,22,24 68:2 <b>everybody</b> 7:9 121:2 134:14 <b>everybody's</b> 80:13 130:21 <b>exactly</b> 8:2 20:2 42:23 48:10 56:18 63:2,15 84:14 107:5 127:14 133:18 <b>examination</b> 5:2 6:13 <b>examined</b> 6:10 <b>example</b> 10:14,15 10:17 12:24 50:20 77:19 81:12 93:8 93:19 100:2 109:15 115:11 119:2 <b>examples</b> 78:21 82:7 <b>excel</b> 94:25 95:5 <b>excellent</b> 117:12 <b>exchanging</b> 95:22 <b>excited</b> 114:17 <b>executes</b> 33:4 60:6 <b>exercise</b> 104:14 <b>exhibit</b> 5:10,11 7:19 9:14 71:1,2,5 <b>exists</b> 15:14 18:1 81:12 <b>exobytes</b> 108:24 119:19 <b>expand</b> 34:21 <b>expect</b> 133:22 <b>expectations</b> 128:5	<b>experience</b> 21:16 34:18,21,23 35:6 37:8 46:20 92:2 93:3,5 94:12,14 96:22 98:10 110:24 <b>experiments</b> 69:23 70:8 <b>expertise</b> 22:7 92:2 102:25 112:17 <b>experts</b> 130:1,17 130:22 131:10 <b>explain</b> 59:19 64:13 87:11 103:15 <b>explained</b> 36:1 <b>explicitly</b> 30:13 46:9 <b>exposed</b> 20:15 24:12 27:11 51:21 51:22 52:25 53:2 53:3 <b>exposure</b> 121:8 <b>extension</b> 30:15 <b>extensive</b> 128:17 <b>extent</b> 89:5 97:23 128:3 <b>external</b> 30:22 32:14 33:19 39:11 48:15 53:16 58:2 58:16 64:20 88:24 95:20,22 113:13 116:15 <b>externally</b> 20:15 24:13 51:9 <b>extra</b> 94:25 <b>extremely</b> 117:13 118:4 120:23 121:21	<b>eyes</b> 1:10 120:14
			<b>f</b> <b>f</b> 50:14 <b>f3</b> 50:4 <b>face</b> 108:1 <b>facebook</b> 1:6 2:6 4:2 6:24,24 7:1,13 10:22,25 12:21 13:6,10 18:10 19:14 20:4,19 22:2,15 23:16,23 24:2 25:4,14 28:4 28:9 30:22 31:19 31:23 33:12,13,20 34:3 35:15 36:11 36:14 37:2,7 40:1 41:16,17 42:8 43:16,16 45:3,4 48:3,7 49:15,24,24 50:3,4,12,13 52:9 53:2,3,18 54:2,3,3 54:10 55:20 56:25 58:4,6,8 59:3,11 61:14 63:24 65:1 65:10,15 67:13,23 68:7 69:10 70:1,6 70:19,21,24 71:25 72:25 73:24 76:11 76:12 77:7,12,23 78:3 80:6 81:6,7 84:8 89:18 90:22 91:6,23 92:23 96:16 97:18 98:4 99:8 100:7,22,25 105:15,25 106:13 106:14 107:21 108:6 110:3,4,11 111:11,12 114:20 117:9,25 120:6 122:19,20,25 123:3,8,18 124:4

**[facebook - friend]**

124:12,20,22	<b>familiarity</b> 12:12	<b>figuring</b> 89:21	<b>forecast</b> 57:10
125:1 128:20	50:18 56:5 63:20	<b>file</b> 71:6 115:14	<b>foregoing</b> 136:7,9
129:3 131:13	66:9 90:24 94:8	<b>files</b> 10:22 95:22	136:14
134:7,24	95:16 99:23 101:1	<b>final</b> 129:10 130:1	<b>forget</b> 20:22
<b>facebook's</b> 18:11	<b>fantastic</b> 84:10	130:11	133:15
35:22 39:13 47:2	<b>far</b> 77:10	<b>financially</b> 136:15	<b>forgets</b> 57:18,18
47:3 61:8 65:21	<b>fashion</b> 81:12	<b>find</b> 12:1 65:23	<b>forgot</b> 120:8
73:10 89:17 100:8	124:6	71:19 115:11	<b>form</b> 10:7 17:3
105:10 107:17	<b>fastdesk</b> 116:10	<b>finds</b> 123:1	82:4,4 96:2 109:5
122:15	<b>faster</b> 19:25	<b>fine</b> 14:5 66:20	112:7 125:3
<b>facilitate</b> 53:19	<b>fastpathdb</b> 51:8	116:17 134:4	<b>formal</b> 35:3,6
<b>facilities</b> 66:15	54:5	<b>finger</b> 19:2	131:15
118:18,19	<b>fb</b> 56:15 60:7	<b>finish</b> 66:13,16,17	<b>forms</b> 81:25
<b>facility</b> 59:22	116:20	83:25	<b>fort</b> 58:15
<b>facing</b> 22:8 105:10	<b>fb.com</b> 4:4	<b>fireside</b> 61:1	<b>forth</b> 14:18 16:23
105:18,20,24	<b>fbcode</b> 54:14 69:3	<b>first</b> 6:9 11:2	36:16 37:25 76:24
<b>fact</b> 26:14	<b>fblearner</b> 55:15	32:12 57:21 81:5	82:10 119:25
<b>fair</b> 53:23 64:2,2	<b>feasible</b> 131:5	89:13,15,21 90:4	120:19 128:18
87:1 99:19 114:6	<b>feature</b> 16:23 29:3	120:25 123:4	130:23 136:8
122:18	29:3,7 50:4,14	124:9	<b>forward</b> 10:6
<b>fairly</b> 10:8 14:7	51:3 64:16 99:1	<b>fit</b> 48:10 107:13	129:15 131:14
69:15	115:3	<b>fits</b> 54:17 64:14	<b>found</b> 120:23
<b>falls</b> 13:4	<b>features</b> 14:12	<b>five</b> 32:2 109:18	<b>four</b> 123:18 131:6
<b>false</b> 105:5 119:17	16:14 50:6 51:5	109:20,21	131:6
126:7 127:11	67:18,20 85:1	<b>flexibility</b> 72:16	<b>framed</b> 124:7
<b>familiar</b> 12:13	94:25 115:3	<b>flow</b> 14:9,11 16:16	<b>framework</b> 16:2
14:16 42:9 44:17	<b>fed</b> 26:4	<b>flowing</b> 41:1	38:2 43:24 44:4
45:12,14 58:20	<b>federal</b> 136:2	<b>focus</b> 128:1	47:21 50:4,14
60:24 61:12,17	<b>federates</b> 115:10	<b>focused</b> 124:6	51:4 59:5 107:8
62:2 63:5 64:6	<b>feed</b> 20:4,6 51:2	<b>fold</b> 19:10	107:12 115:9
65:4,17 66:3	76:19,23,23 81:13	<b>folks</b> 130:7	127:25
68:19 69:13,15,21	81:14 90:8,9,9,9	<b>follow</b> 29:24 71:18	<b>framing</b> 118:15
72:9,11 73:15	90:10,13 113:3,7	73:23 83:23 84:24	<b>frankly</b> 121:18
77:6 84:9 85:5	<b>feel</b> 117:10,25	118:8 120:7	128:23
86:13 91:13 95:12	130:25	121:15 126:19,24	<b>free</b> 117:10,25
97:10,15,17,22	<b>fi</b> 111:9,13	128:16 132:7	130:25
98:15,18 99:12,21	<b>figure</b> 17:19 47:11	<b>following</b> 130:22	<b>fresh</b> 30:10
103:11,14 106:7	80:18 86:4,5	<b>follows</b> 6:11	<b>friday</b> 1:15 2:13
111:7 114:14,21	89:13 96:7 108:19	<b>fonti</b> 3:4	6:1 121:13
116:6	<b>figures</b> 28:20	<b>foqs</b> 57:2	<b>friend</b> 20:20 21:21
			90:9

[frl - help]

<b>frl</b> 58:15 <b>front</b> 7:3 119:13 130:15,18 <b>fs</b> 42:16,17,19 <b>fti</b> 133:13 <b>fully</b> 57:14 86:3 134:21 <b>fulmer</b> 1:18 2:11 6:6 136:24 <b>function</b> 22:13 63:20 89:23,24 90:2 125:4 <b>functional</b> 33:11 <b>functionality</b> 51:25 <b>further</b> 17:22 123:19 124:11 136:13,15	<b>gateway</b> 70:12 <b>gathered</b> 104:15 <b>gathering</b> 104:23 <b>gauge</b> 125:23 <b>gdpr</b> 15:23 <b>general</b> 38:25 53:14 69:19 86:3 87:25 112:3 <b>generally</b> 78:7 <b>generate</b> 36:16 <b>generated</b> 56:12 65:9 82:10 <b>generates</b> 20:6 <b>generation</b> 97:9 110:22 <b>generic</b> 72:6 <b>genomized</b> 81:16 <b>genre</b> 40:10 <b>geographically</b> 19:21 <b>getting</b> 10:14 49:8 103:1 <b>gibberish</b> 83:20 <b>gibson</b> 3:17,21 <b>gibsondunn.com</b> 3:19,24 <b>give</b> 10:15 12:24 17:24 29:11 50:13 82:24 91:23 98:20 113:20 119:21 123:19 <b>given</b> 8:25 11:14 61:4 78:15 100:3 <b>glad</b> 117:15 <b>gluster</b> 42:16,17 42:19 <b>go</b> 9:11,14 11:17 11:18 14:7 15:6 16:4 18:25 20:4 20:12 24:18 27:6 30:3,13 34:7	38:18,20 51:19 66:15 68:10 71:17 71:22 73:18 77:7 77:23 79:16 84:6 87:22 109:3 112:5 112:6,6 117:3,5,9 120:3,3,25 121:6,9 121:13 123:19 124:10 125:19,24 128:21 129:11 132:11,12 135:2 <b>goal</b> 126:11 <b>goals</b> 59:8,9,10,12 <b>goes</b> 13:4,15 43:17 45:9 56:22 82:19 <b>going</b> 6:4 11:19 12:24 20:12 24:21 38:19 41:21 42:2 44:22 46:15 51:19 51:20 59:16 66:11 83:11 84:4 86:19 90:2,6 92:13 109:7 121:6 123:15,21 125:5 127:18 133:14 134:1 <b>good</b> 6:15,15 7:7 15:21 26:21 31:9 84:21 116:1 121:23 123:12 <b>google</b> 92:14 <b>googled</b> 47:11 <b>grant</b> 113:25 <b>granted</b> 114:5 <b>graph</b> 11:4 28:25 46:16,17 101:3 102:11,15,17,19 112:17,23,25 113:3,7,8,20 <b>graphics</b> 22:14	<b>graphing</b> 102:23 <b>grateful</b> 121:20 130:3 <b>great</b> 31:6 78:13 122:2 <b>green</b> 85:11,14 <b>groups</b> 50:22 107:21,23 <b>guess</b> 9:12 14:6 17:12 22:4 23:11 25:25 26:9 55:8 75:7 99:15 111:1 112:1 134:13 <b>guys</b> 29:21 34:4 55:13
<b>g</b>	<b>g</b>	<b>h</b>	<b>h</b> 5:8 <b>hacker</b> 4:3 <b>half</b> 122:8 <b>hand</b> 17:25 <b>handful</b> 128:17 <b>handled</b> 30:5 <b>happen</b> 28:14 107:11 <b>happened</b> 12:25 <b>happy</b> 66:23 <b>hardware</b> 38:3 <b>hash</b> 83:19 <b>hashed</b> 10:22 11:1 <b>hat</b> 116:19 <b>headspace</b> 131:9 <b>hear</b> 103:16 118:9 <b>heard</b> 26:11 59:18 123:15 134:5 <b>hearing</b> 1:13 2:10 117:20 122:8 129:13,15,18 131:10 135:5 <b>help</b> 82:8 84:7,10 117:15 125:9



[helpful - interested]

<b>helpful</b> 47:10 117:13 118:4 120:24 121:22 <b>hexadecimal</b> 103:17 <b>high</b> 7:24 8:6 10:9 35:24 <b>higher</b> 19:4 <b>highly</b> 1:9 65:23 115:19 120:13 <b>hints</b> 16:22 <b>historical</b> 54:22 74:23 <b>history</b> 50:10 <b>hive</b> 59:18,18,19 59:23 60:4,7,22 119:3,4 <b>hiveql</b> 60:7 <b>home</b> 92:21 93:4 <b>homestretch</b> 116:9 <b>honest</b> 73:16 <b>honestly</b> 25:20 <b>honor</b> 120:22 <b>hope</b> 114:8 <b>hopefully</b> 66:17 81:8 126:15 129:14 <b>hoping</b> 48:11 117:5 <b>hotline</b> 61:1 <b>hour</b> 66:18 <b>hours</b> 52:18 122:8 <b>house</b> 84:11 <b>hr</b> 59:8 <b>http</b> 67:5 <b>huge</b> 119:19 <b>huh</b> 7:20 10:20 22:16 47:14 48:25 55:16 62:5 69:16 72:20 73:20,20 78:21 79:21 90:21	92:22 100:20 107:9 108:4 <b>hulkdb</b> 61:11,15 114:11 <b>human</b> 81:8 103:15,16,20,23 104:5,11,19,25 105:8 <b>hundred</b> 128:13 <b>hundreds</b> 115:20 <b>hyper</b> 67:6 <b>hypothetical</b> 124:2	<b>ignorant</b> 34:13 <b>illustrate</b> 13:1 <b>illustrative</b> 12:24 <b>image</b> 20:7 22:20 <b>imagery</b> 116:3 <b>images</b> 19:19 <b>imagine</b> 30:11 47:4 108:21 120:8 <b>impasse</b> 134:8,11 <b>import</b> 94:23 <b>importantly</b> 46:21 91:11 <b>impressions</b> 74:10 <b>improve</b> 105:9 <b>improved</b> 111:3 <b>inbound</b> 124:21 125:5 <b>incoming</b> 96:11,15 <b>indexing</b> 115:1 <b>indicate</b> 26:15 36:22 <b>indicated</b> 15:17 96:2 <b>indicates</b> 11:21 37:3 133:6 <b>indicating</b> 38:24 <b>indication</b> 36:20 130:24 <b>indicator</b> 48:23 <b>individual</b> 44:22 53:20 78:8 79:19 79:24 80:8 82:20 82:20 83:17 <b>individuals</b> 56:25 <b>infer</b> 16:24 82:12 82:18 <b>inference</b> 38:8 <b>inferences</b> 111:25 <b>inferring</b> 46:19 76:22	<b>information</b> 9:20 10:10 13:25 15:4 17:22 20:24 21:1 21:2 28:10 29:9 32:1,4 53:7 59:11 71:19 79:19 84:17 114:3 118:22,23 120:15 125:22 127:4,19,22 129:3 129:7 <b>informative</b> 81:20 121:21 <b>infrastructure</b> 46:16 50:3 58:4,7 58:9 77:4 87:17 113:14 <b>initial</b> 23:5 87:15 <b>initiated</b> 32:6 <b>initiates</b> 14:1 <b>initiatives</b> 58:16 <b>inner</b> 97:8 <b>insights</b> 44:15 45:2,3 <b>instagram</b> 114:12 114:18,22 <b>instance</b> 42:19 61:22 96:10 108:15 113:2,6,12 119:21 <b>instances</b> 69:20 110:19 111:19 <b>instantaneous</b> 52:18 <b>intake</b> 10:7 <b>integrated</b> 47:12 <b>intelligence</b> 55:18 <b>intends</b> 122:19 <b>interactive</b> 46:24 <b>interested</b> 16:7 23:10 26:8 62:25 132:2 136:16
	<b>i</b>		
	<b>ian</b> 4:2 <b>ianchen0</b> 4:4 <b>ibubiz</b> 62:14 <b>idea</b> 13:1,14 14:10 37:17 40:12 45:5 45:6 68:18,25 72:8,10 73:5 77:18 85:2 87:25 99:4 126:15 129:18 <b>ideally</b> 71:8 <b>ideas</b> 69:12 86:7 <b>identifiable</b> 79:11 79:19,23,23 80:5,6 80:7 <b>identified</b> 5:10 11:11 24:25 33:9 127:10,12 <b>identify</b> 8:7 11:22 12:2 15:13 83:18 105:14 123:5 126:2 128:4 <b>identifying</b> 15:15 <b>ids</b> 10:25 56:9 83:18 <b>igml</b> 114:25		

**[interests - know]**

<b>interests</b> 82:13 105:9 <b>interface</b> 51:15,21 93:17 110:8 113:17 <b>interfaces</b> 53:14 <b>intermediary</b> 110:8 <b>internal</b> 14:13 16:17 17:9 18:11 19:18 20:5 23:18 23:20 24:2,7 25:3 25:5,7,10,13,25 30:21 31:4,8,23 35:11,13 36:4,5,8 36:10 37:24,25 38:21 39:11,24 40:4,5,8,12 41:3,8 42:12,14,20,24 44:11,13 45:8,9,15 46:10,11 47:15,16 48:7,19,25 49:10 49:12,23 50:1,11 51:6,7 52:23 53:18,25 54:1,2 55:6 57:9 58:6 59:8,13,24 61:22 61:23 62:10,18,23 63:7 66:1 68:5 69:24 70:6 74:18 76:16 77:4 84:8 85:16 89:17 92:12 94:15,18 96:25 97:24,25 98:3,11 101:5 102:3 103:3 113:16 114:11 115:17 116:12,13 <b>internally</b> 31:19 33:9,13 34:4 35:4 41:24,25 45:24 46:1 49:25 50:19	56:11 60:23 65:1 81:14 85:20 101:21 <b>interpret</b> 16:13 65:25 <b>interpretation</b> 74:15 <b>interview</b> 30:4 <b>inventory</b> 7:13 25:24 53:11 <b>investigate</b> 46:5 <b>investigation</b> 45:17,20 <b>involve</b> 50:14 <b>involved</b> 30:9 53:9 66:6 68:24 69:1 76:19 116:4 <b>iqr</b> 33:19 <b>irc</b> 61:21 <b>issue</b> 33:11 122:13 124:8 129:12,13 130:1,11,25 131:1 <b>issued</b> 131:6 <b>issues</b> 134:7,12 <b>issuing</b> 131:15 <b>iterative</b> 124:6  <b>j</b> <b>january</b> 1:15 2:13 6:1 136:19 <b>job</b> 25:1,10 33:22 <b>jobs</b> 32:23 33:14 96:23 98:1 <b>joe</b> 78:9 <b>join</b> 16:11,12 50:22 117:25 <b>joins</b> 95:2 <b>jsc</b> 1:6 2:6 <b>judge</b> 124:7 125:7 126:20 128:25 129:4 134:9,19	<b>judge's</b> 122:6 128:22 129:16 <b>jump</b> 117:17 <b>jupyter</b> 58:21,23 58:24,25 59:2 <b>jupyterhub</b> 58:25 <b>jurat</b> 136:2  <b>k</b>  <b>k</b> 3:5 <b>keep</b> 20:12 51:19 <b>keeping</b> 41:1 125:4 <b>keller</b> 3:10 <b>kellerrohrback.c...</b> 3:13,14 <b>kept</b> 125:5 <b>key</b> 39:1 82:19 83:21 88:23 112:5 112:11 <b>kid's</b> 49:6 <b>kind</b> 8:23 9:17 13:20 14:3,16 15:12 16:15 22:6 32:10 33:24 34:19 34:25 36:15 38:25 41:8 44:21,23 46:19 52:11 64:15 67:19 70:3 72:5 76:8,14 78:11 82:22 83:6 86:19 91:5,21 92:1 96:9 96:24 97:23 102:9 102:13,21 103:24 104:18 109:23 110:8,17 111:2,3 112:4,10 115:15 115:22 116:14 131:9 <b>kinda</b> 10:7 19:10 31:16 40:4 47:11 87:13 93:7,24	94:22,24 95:2 <b>kinds</b> 118:21 <b>knew</b> 8:14 11:23 <b>know</b> 8:14,16,23 9:3,23 10:1,3,5,11 11:2,3,7,9,19 12:4 12:7,9,11 13:10,19 13:19 14:3,6,8,15 14:19,23,24 15:8,9 15:13,14 16:5,20 16:21 17:11,25 18:2 19:20,23 20:3,8,15,21 21:5 21:5,6 22:4 23:4,9 23:14,17,22 24:14 24:17,24 25:20 26:24 27:3,13,13 28:9,16 29:9,14,17 30:23,24 31:24 32:8,15,24 33:24 34:13,15,20,25 35:4,23 36:24 37:19,20 38:6,10 38:13,14,22 39:4 39:15 41:13 43:1 43:16 44:15,21 45:1,4,21 46:2,8 46:11,18 47:4 48:13,15 49:16,17 50:10,20 51:14 52:20,21 53:9 54:20 55:4,23 56:9,20 58:17,19 58:21,22 59:11,12 59:14,15,22 60:8 60:13,15,16 61:7 61:10,17,19 62:5,9 62:12,15 63:5,9,21 63:22 64:2,3,9,10 64:18 65:8,20 66:11 67:6,12,13
--	--	--	--

[know - logging]

67:23 68:22 69:6 70:10,15,16,19,22 72:2,12,14,21 73:8 73:16,21,22 74:2 74:25 75:1,2,7,14 75:20,23,24 76:1,4 76:9,15,22,23,24 77:1,15,20,21 78:4 78:17 79:5,7 80:1 82:2,4,7,9 83:17 84:16,16 85:12,15 85:17,24 86:1,22 87:6,10,14,19,20 88:9,12,13,16,17 88:20 89:5,8 90:3 90:14,17 91:5,7,9 91:21,21 92:18 93:6,13 94:3,22,24 95:2 96:1,2,18 97:5,7,8 98:7,22 99:5 100:5,7,10,10 100:22 101:4,9,14 101:18 102:13,14 102:16,24 103:23 103:24,24 104:16 104:16,18,21,25 105:12 106:14,18 106:18 108:14,23 109:2,4,16,20,23 109:24 110:16 111:2,12,14,20 112:18 113:6 114:1,7,19 115:2,5 115:12,12,16,20 116:1,14,15 118:22 119:2,11 121:8 123:18 125:8,15 126:7 127:24 128:2,5 132:17,18,24	<b>knowledge</b> 12:15 13:2 16:25 17:12 18:1,20 22:7 24:22 25:17 27:22 29:7 31:1,16 32:21 40:3,13 43:6 46:20 47:22 52:5 54:12 55:16 57:16 59:2 60:5 60:12 73:4,7 86:9 87:19 88:22 91:22 92:10,11 94:11 107:25 111:10 118:14 119:23 <b>known</b> 108:16 <b>kutscher</b> 3:17 80:12 84:9 133:12 <b>l</b> <b>l</b> 42:4 111:6 <b>lab</b> 47:5 <b>labs</b> 72:4 91:2 <b>lakis</b> 61:24 <b>language</b> 60:8 103:23 104:17 <b>laptop</b> 53:13 <b>large</b> 46:17 60:13 102:22 124:1 <b>larger</b> 47:12 <b>laser</b> 62:4,5 <b>laufenberg</b> 3:11 <b>laura</b> 3:22 <b>lawyers</b> 117:11,19 121:7 <b>layer</b> 23:23 24:3 67:19 93:20 108:16 <b>ibu</b> 62:10 <b>lead</b> 30:18,18 <b>leaf</b> 63:6 86:11,11 86:14	<b>leaf4</b> 62:24 <b>leaf5</b> 63:7,9 <b>learn</b> 119:9 <b>learned</b> 11:14 <b>learning</b> 8:13 50:6 50:9,11,19,21 51:5 51:5 55:18 60:14 114:12 <b>leave</b> 46:12 49:14 81:7 120:17 <b>leaving</b> 117:22 <b>legal</b> 4:2 29:23 63:18 84:5 <b>lesley</b> 3:4 <b>letter</b> 10:19 70:25 71:2 <b>letters</b> 83:20 <b>level</b> 7:24 8:6 10:9 32:2 33:12 70:3 70:10 73:7 78:11 80:8 87:12 88:1 88:12 92:17 93:17 100:11 102:3 103:8 106:16 114:3 116:2 127:8 127:20 <b>libraries</b> 93:25 111:4 <b>library</b> 97:24 <b>lidbpostmortem</b> 64:25 <b>life</b> 14:3 59:9 <b>lift</b> 74:1 <b>light</b> 134:18 <b>likes</b> 100:4 <b>limit</b> 15:22 73:7 126:1 <b>limitation</b> 123:1 130:16 <b>limited</b> 15:13 55:20 107:20	122:23 <b>limits</b> 118:14 <b>line</b> 52:22 55:12 59:22 60:13 71:17 72:6 81:10 83:25 128:6 133:7 <b>lines</b> 32:2,3 86:23 <b>linkable</b> 80:20 <b>linked</b> 80:23 <b>linking</b> 102:14 <b>list</b> 7:18 11:25 13:18 14:7 15:18 16:4 19:11 26:14 118:20 125:10 132:16 <b>listen</b> 118:11 <b>litigation</b> 1:7 2:7 <b>little</b> 12:15 19:4 23:17 24:22 27:22 31:1 32:20 43:6 52:5 56:5 58:12 59:19 79:25 83:8 131:20 <b>liv3</b> 64:5 <b>live</b> 64:4,7,12,14 64:15,18 69:12 109:25 <b>llp</b> 3:10 <b>lmmum</b> 3:24 <b>load</b> 33:25 <b>local</b> 36:3 <b>locality</b> 27:21 <b>location</b> 82:11 <b>loeser</b> 3:11 <b>log</b> 91:4,6 <b>logarithm</b> 65:6 <b>logdevice</b> 65:25 66:4 <b>logged</b> 65:12 <b>logging</b> 65:9 91:7 125:4
--	--	--	--

[logs - mobile]

<b>logs</b> 89:22,23,25 <b>long</b> 35:7 66:11 121:13 <b>longer</b> 99:18 <b>look</b> 7:8 9:13 10:21 18:16 28:13 129:11 <b>looked</b> 57:13 60:9 90:19 <b>looking</b> 7:10 16:2 55:3 71:10 80:10 80:14 81:21 123:13 125:22 <b>looks</b> 23:2 77:3 89:9 <b>lose</b> 39:16 <b>losing</b> 54:25 <b>lost</b> 133:10 <b>lot</b> 9:14 26:11 39:6 55:16 84:12 103:16 118:6 120:15 121:19 124:10,11 129:6 <b>lots</b> 59:18 124:14 <b>love</b> 66:21 <b>loves</b> 121:7 <b>lower</b> 92:17 93:17 93:20 102:3 106:16 <b>lucky</b> 122:7,9 <b>lweaver</b> 3:7	<b>maintain</b> 9:25 <b>maintained</b> 118:17 <b>making</b> 50:21 63:24,25 70:19 72:22 84:4 89:23 89:24 107:17 117:3 121:21 <b>manage</b> 14:14 15:24 32:24 34:4 34:9 53:22 66:5 <b>managed</b> 8:15 86:3 115:18 <b>management</b> 39:9 39:22 41:9 52:23 116:20 <b>manager</b> 45:7 <b>manages</b> 11:24 108:15 <b>managing</b> 51:4 66:1 68:16 <b>manifold</b> 66:25 67:4 <b>manipulations</b> 95:3 <b>mann</b> 4:7 <b>map</b> 15:23 116:3 <b>mapping</b> 116:5 <b>mappings</b> 56:8 <b>marked</b> 40:22 120:13 <b>martie</b> 3:17 118:3 <b>massive</b> 124:14 <b>master</b> 1:11 133:10 <b>matches</b> 11:1 <b>matching</b> 10:24 <b>matter</b> 122:21 <b>matthew</b> 3:5 <b>md</b> 1:6 2:6	<b>mdl</b> 1:5 2:5 <b>mean</b> 12:1,3,5 13:12 14:8 16:12 17:10,11 18:18 28:3,4 30:2 34:12 36:1 53:3 54:1,2 61:2,2,5 69:18 77:18 79:18 83:10 83:18 85:17 86:7 88:7,10 90:23 93:15 94:8 97:19 97:22 98:17 99:18 103:21 104:1 110:23 114:8 118:12 119:19 132:5 <b>meaning</b> 54:2 65:22 79:11,23 <b>meaningful</b> 105:2 <b>means</b> 14:10,19 40:17 44:19 53:1 54:23 65:1 67:7 104:2 <b>meant</b> 44:16 85:12 85:14 86:5 97:9 <b>media</b> 19:14 <b>mediation</b> 115:23 <b>meet</b> 9:6 13:22 29:16 109:8 <b>meeting</b> 29:24 <b>meets</b> 80:3 <b>melamed</b> 3:5 <b>memcache</b> 68:4 <b>memory</b> 30:10 32:12 40:11 63:11 103:5 <b>menlo</b> 4:3 <b>mentioned</b> 17:6 <b>mergers</b> 36:23 <b>message</b> 51:12 52:11 96:23	<b>messages</b> 23:22 51:13 52:10 72:6 <b>messaging</b> 23:21 24:12,13,17 51:21 51:25 53:2,3,19,19 54:3,4,6,10 96:24 <b>messenger</b> 51:13 52:6,7,9 53:12,17 73:17 <b>meta</b> 28:7 33:8 93:22 94:20 97:25 98:7 113:16,16,19 115:25 <b>meta's</b> 113:14 <b>metadata</b> 62:25 68:16 <b>metastore</b> 68:16 <b>method</b> 95:21 <b>methodical</b> 131:12 <b>metric</b> 77:17 83:15 <b>michael</b> 4:7 <b>michelle</b> 1:18 2:10 136:24 <b>migrate</b> 23:13 <b>milan</b> 1:18 2:10 6:6 136:24 <b>mill</b> 3:18 <b>minute</b> 117:4,6 120:2 <b>minutes</b> 66:14 <b>misrepresenting</b> 53:24 58:1 <b>misstep</b> 84:5 <b>misunderstood</b> 133:20 <b>mkutscherclark</b> 3:19 <b>mmelamed</b> 3:8 <b>mobile</b> 37:8 58:6,8
<b>m</b>			
<b>m</b> 36:18,22 46:25 47:6 <b>machine</b> 50:5,9,11 50:19,20 51:5,5 55:17 60:14 86:12 104:1,2,6,18 114:12 136:12 <b>machines</b> 86:14			

**[mode - operates]**

<b>mode</b> 34:19 <b>models</b> 47:21 <b>monitor</b> 125:2 <b>mononoke</b> 68:20 68:21 <b>morning</b> 6:15 122:8 134:11 <b>mother</b> 28:9 <b>motion</b> 134:1,6,16 <b>mouthful</b> 115:4 <b>move</b> 28:20 35:2,5 129:15 131:7 <b>moved</b> 77:21 98:23 134:9 <b>movement</b> 83:15 <b>movements</b> 77:17 <b>moves</b> 28:19 <b>moving</b> 28:22 <b>muddler</b> 114:11 <b>multi</b> 69:12 <b>multiplay</b> 35:10 <b>multiple</b> 69:20 82:3 <b>mumm</b> 3:22 126:23 127:21 130:14 131:19,23 132:7 133:4,5 <b>music</b> 37:15 <b>mysql</b> 17:9 23:15 69:15 75:16,16 96:10 102:12 108:12 110:17,19	128:3 129:1 <b>narrowed</b> 125:7 <b>nature</b> 10:2,8 96:3 <b>nearly</b> 122:9 <b>necessarily</b> 26:19 26:23 32:7 80:3 <b>necessary</b> 15:3 126:18 <b>need</b> 25:20 31:8 56:19 84:17 88:13 89:13 116:1 120:11 123:19 125:19,21 128:24 129:10 <b>needs</b> 40:21 109:16 126:8 <b>neither</b> 136:15 <b>network</b> 13:7,11 13:12 14:14 16:18 18:23 20:16 49:2 82:12 111:12 <b>networks</b> 123:25 <b>never</b> 45:9 56:22 56:24 81:15 <b>new</b> 3:23,23 23:2 34:18,19,21,23 35:5 57:10 69:23 69:24 70:5,8 <b>news</b> 20:4,5 76:19 76:22,23 81:12,13 90:13 <b>nine</b> 122:13,17,17 <b>non</b> 28:7 105:15 <b>northern</b> 1:2 2:2 <b>note</b> 133:5 <b>notebooks</b> 59:7 <b>noted</b> 120:16 123:2 130:12 133:2,9 <b>notes</b> 103:7	<b>notification</b> 56:10 <b>notifications</b> 56:7 56:12 <b>novi</b> 47:24 48:13 48:16,22 70:12 97:2 <b>npe</b> 34:18 57:11 70:2 <b>number</b> 60:14 100:4 <b>numbers</b> 83:20	35:18,20 36:1,3 39:4,9,23 40:23 43:21 45:16 47:10 47:13,21 48:2,20 48:24 49:3,17 51:7,11 53:18 54:9,14 55:15 56:14 58:12 59:8 59:16 60:16 62:13 62:22 63:7,18 64:7,12 65:25 66:10 68:10,15,20 70:12,16,24 71:12 72:5,10 73:5 83:1 85:9,24 86:3,11 87:1 90:7 93:14 93:23 95:18 96:19 97:8,17,19 99:2 106:2 107:20 108:3 109:7 114:5 114:16 116:1 117:1,8,22 120:23 122:4,6 130:12 132:4 134:3 <b>old</b> 40:16 52:17,18 109:22 <b>older</b> 132:23 <b>omicron</b> 49:7 <b>omnistore</b> 73:5 <b>once</b> 128:11 131:13,14 <b>ones</b> 16:5 51:15 123:24,24 <b>online</b> 30:20 <b>open</b> 46:22 118:6 123:4 <b>operate</b> 50:9 <b>operated</b> 8:12 <b>operates</b> 21:6 49:1 92:17
<b>n</b>		<b>o</b>	
<b>n</b> 5:1 <b>name</b> 36:3 71:6 74:4 81:5,6 133:16 136:19 <b>named</b> 122:16,17 122:17,24 123:7 <b>narrow</b> 122:12 125:6 126:11		<b>oakland</b> 3:6 <b>oath</b> 136:10 <b>obfuscate</b> 55:1 <b>objects</b> 63:12,13 <b>obs</b> 54:24 <b>obscuring</b> 54:24 <b>observe</b> 82:11 <b>obvious</b> 95:7 119:8 <b>obviously</b> 73:18 117:8 125:10 <b>occurring</b> 83:10 <b>oculus</b> 71:24,25 72:4 <b>offd</b> 72:5 <b>offend</b> 17:16 <b>offer</b> 48:13 <b>offering</b> 13:9 <b>oh</b> 28:20 35:20 40:21 49:9 58:12 62:13 <b>oil</b> 72:10 <b>okay</b> 6:15 7:1,7,18 7:21,24 11:10 13:1,24 14:5,24 15:18 17:1,13 19:9 22:9 27:19 30:14,15,20,24,25 31:6,10,18 34:12	



**[operational - pipelines]**

<b>operational</b> 21:15 <b>operations</b> 25:3 40:1 62:11 66:1 <b>opinion</b> 92:2 <b>opportunities</b> 119:9 <b>opportunity</b> 73:25 129:24 130:4,9 134:17 <b>opposed</b> 131:10 <b>ops</b> 98:4 <b>optimize</b> 16:18 <b>option</b> 33:21 <b>options</b> 10:1 89:3 <b>order</b> 13:23 23:13 87:22 122:22 123:13 129:4,14 130:15 131:3,16 134:19 <b>orderdb</b> 74:23 <b>ordered</b> 129:1 <b>orders</b> 122:7 131:6 <b>organization</b> 116:25 <b>oriented</b> 32:5 <b>outbound</b> 125:6 <b>output</b> 13:6 49:13 49:14 50:24 51:1 <b>outputs</b> 13:5,18 <b>outside</b> 13:20 22:6 28:7,7 35:4 51:21 51:23 53:1 70:20 84:10 90:6 92:1 112:16 113:8,14 <b>owned</b> 116:22 <b>owner</b> 115:24	<b>pacific</b> 2:13 <b>package</b> 56:15 116:19,20 <b>page</b> 3:18 5:4,9 19:10 20:6 66:13 122:14 <b>palo</b> 3:18 <b>panelist</b> 10:25 <b>paragraph</b> 7:12 <b>parallel</b> 32:22 <b>parens</b> 11:20 38:23 <b>park</b> 3:22 <b>part</b> 16:15 18:2 23:10 29:14 30:12 30:13 37:1,1,4 47:25 53:25 55:17 57:3 62:24 72:3 78:11 82:19 96:24 99:17 102:22 104:19 120:3 121:1 123:10 127:11 <b>particular</b> 9:7 12:23 18:19 30:5 30:8 57:17 81:13 90:1,2 118:7 121:16 124:22 <b>particularly</b> 8:6 62:24 123:24 <b>parties</b> 12:5,8 18:11 21:24 26:7 27:11,12 32:14,14 35:13 37:18 40:25 41:4,14 42:7 46:22 47:18 50:15 57:7 59:25 60:2 62:8 63:4 64:1,8 65:3,22 66:15 67:24 69:17 72:13 73:13 75:10,22	85:21 87:4 89:17 89:19 90:5,15 91:12,24 94:4 95:8 96:21 97:6 97:18 98:21 111:19 112:14 113:5 114:13 116:4 118:25 119:4 120:17 122:25 130:22 136:17 <b>partner's</b> 10:25 <b>parts</b> 12:21 107:13 <b>party</b> 13:10 21:20 24:4 53:4,5 56:1 60:3 64:20 65:15 68:7 69:10,12 70:7 72:22 74:20 77:11,14 85:8 89:23 92:24 93:11 99:9 100:6 107:17 110:11 <b>passed</b> 18:25 <b>path</b> 87:13 <b>patterns</b> 77:25 <b>pausing</b> 82:23 <b>payer</b> 39:1 <b>payment</b> 75:1,4 91:24 <b>payments</b> 75:15 <b>pdf</b> 7:21 <b>peak</b> 32:24 <b>pending</b> 66:17 <b>people</b> 41:12,13 48:15 50:16 59:9 75:11 77:21,23 78:7,20 79:13 80:22 82:18 84:12 91:17 92:24,25 105:17 125:20 128:15,17	<b>perfect</b> 25:21 <b>perfectly</b> 73:23 <b>perforce</b> 76:2,3,4 76:5,6,7,10,12,15 <b>performance</b> 19:13 33:25 40:7 68:6 <b>performing</b> 99:1 <b>performs</b> 16:11 <b>period</b> 132:18 <b>periplus</b> 76:18 <b>persistent</b> 52:3 <b>person</b> 26:1,9 28:13 79:20,24 80:5,20 88:11 105:20 119:14 <b>personal</b> 13:2 <b>personalization</b> 105:9 <b>perspective</b> 59:20 <b>phase</b> 110:25 <b>phone</b> 53:13 56:10 56:11 <b>phonetic</b> 133:13 <b>photo</b> 19:17 83:5 <b>photos</b> 21:4 22:14 49:21 63:13 <b>phrase</b> 86:10 <b>phrased</b> 107:6 108:11 110:6 <b>physical</b> 62:20 63:13 92:19 <b>picture</b> 20:5,13,18 20:21 21:10,22 <b>pictures</b> 21:3 <b>piece</b> 27:4 38:4 57:14 104:12 <b>pii</b> 10:23 11:1 <b>pinpoint</b> 77:3 82:1 <b>pipelines</b> 111:5
<p><b>p</b></p>			
<b>p</b> 42:4 <b>p.m.</b> 2:12 6:2,2 68:13,14 135:5			

[place - provide]

<b>place</b> 104:24 136:8 <b>placed</b> 136:10 <b>places</b> 39:6 <b>plaintiff's</b> 122:24 <b>plaintiffs</b> 3:3 118:7 122:16,17 123:7 129:24 132:14,24 134:5,7 134:13 <b>plan</b> 23:1,6 131:5 <b>platform</b> 13:7,9,15 15:7 16:8,9 18:10 21:20 22:15 23:8 23:16 33:2,5 35:15,22 37:2,18 39:13,14 42:5 44:10 45:9 47:2,3 47:12 48:3,21 50:16 55:8,21,22 56:21 57:8,22 61:8 65:21 67:25 72:1 75:11 89:17 90:15 96:16 100:8 100:25 106:14 108:22,22 116:3 <b>play</b> 19:15 37:21 <b>playing</b> 91:17 <b>please</b> 16:6 <b>pleasure</b> 122:1 <b>plus</b> 40:17 <b>point</b> 8:20 14:25 15:15 17:5 22:20 56:21 76:11 79:10 83:21 87:8,15,20 87:25 99:2 104:13 106:24 120:1 122:11 125:18 131:14 134:11 <b>pointed</b> 86:20 <b>policies</b> 54:20 89:4	<b>policy</b> 9:8 27:4 109:9 <b>pool</b> 83:17 <b>pope</b> 5:3 6:8,16 66:19 68:18 69:12 72:8,10 80:15 117:5,9,12 119:10 119:12,13 120:3,7 120:25 121:5,12 125:14 128:4 130:7 131:16 <b>pope's</b> 118:13 <b>portion</b> 102:11 112:18 113:13 122:15 <b>poseidon</b> 84:25,25 <b>position</b> 27:2 <b>possibility</b> 8:18 <b>possible</b> 23:7 26:17 27:1 30:7 52:1 66:22 69:5 86:22,23 127:15 <b>post</b> 100:4 133:12 <b>posted</b> 20:5 134:10 <b>posts</b> 63:14 76:24 82:10 <b>potentially</b> 115:21 117:19 <b>power</b> 85:15 112:22 <b>powers</b> 37:13 <b>practices</b> 54:22 <b>preparation</b> 16:16 <b>prepared</b> 132:16 132:19 <b>presence</b> 85:11 <b>present</b> 4:6 <b>presto</b> 23:7 <b>pretty</b> 23:14 25:18 87:14 95:7 101:12	128:23 129:4 <b>previous</b> 108:10 <b>previously</b> 107:10 112:3 123:2 <b>primarily</b> 49:22 <b>primary</b> 102:21 <b>prior</b> 136:10 <b>privacy</b> 1:6 2:6 15:24 22:18 115:5 115:8,8,9,22 <b>privilege</b> 133:17 <b>probably</b> 13:13 18:20 29:3,22 34:20 57:16 66:17 105:16 112:8 118:19 123:17 124:20 <b>problem</b> 23:1 72:19 134:3 <b>procedure</b> 130:2 <b>proceed</b> 129:8 <b>proceeding</b> 71:1 <b>proceedings</b> 136:7 136:9,11 <b>process</b> 9:22 10:24 29:20 30:12 36:25 49:13 84:3,8,14 89:1 97:1 121:18 128:10 133:16,24 <b>processed</b> 13:16 14:2 15:11 32:7 <b>processing</b> 36:15 64:12,18 85:25 86:4 96:23,25 97:24 98:1 <b>prod</b> 28:16 <b>produce</b> 133:7 <b>producing</b> 133:14 <b>product</b> 12:22 26:18 28:8,16 29:2 33:20 34:18	34:19,21,23 35:5 43:17 53:9 57:10 64:16 69:23,24 70:5,8 72:4 88:25 93:7 97:8 106:7 108:11 <b>product's</b> 87:23 <b>production</b> 35:3 46:3 94:21 133:18 133:23 <b>productive</b> 42:2 <b>products</b> 19:14 24:7 28:2 52:9 64:17 68:6 73:11 77:25 79:6,8 80:3 82:5 86:24 87:17 101:22 105:10,11 105:15,18,19,20 105:24 112:9 115:13 118:24 <b>profile</b> 1:6 2:6 91:5 101:5 <b>program</b> 8:12 10:5 13:21 53:10 115:9 <b>progress</b> 115:19 <b>prometheus</b> 23:13 <b>pronounce</b> 41:19 62:15 <b>pronunciations</b> 37:22 <b>protected</b> 52:10 <b>protective</b> 130:15 <b>protocol</b> 84:6 133:6 <b>prototype</b> 34:19 <b>provide</b> 29:5,10 54:23 78:23 84:3 87:24,25 95:18 116:15 123:14 126:16 129:3
--	---	--	---

[provided - referenced]

<b>provided</b> 7:18 23:11 78:21 82:9 104:10 127:19	<b>question</b> 7:24 11:2 11:8 12:17 20:21 32:13 38:15 42:6 49:11 51:17 55:8 57:5,21 67:2 73:12 80:16 82:24 83:7,23 87:3,6 88:14 90:4 100:22 101:3 106:11 108:6,17 109:10 109:24 111:17 112:24 119:6,15 119:17 122:10 123:4,5 124:19 125:15,16 127:3 128:22,23 132:13 132:20 133:11	<b>quickly</b> 14:7 27:7 125:12,15	112:7 117:13 121:20
<b>provides</b> 13:7 56:8 67:17 72:15 88:22 112:25 113:17	<b>questioning</b> 131:24	<b>r</b>	<b>reason</b> 18:6 24:25 41:11 48:5 55:19 67:4 72:15 73:17 77:10,12,22 83:14 84:22 99:7,16
<b>providing</b> 29:8 31:7 51:4 92:18 96:9	<b>questionnaire</b> 26:17 29:15,21 80:9 92:5 105:23 126:9,10,12,16 127:2,6,13 131:17 131:17,21 132:2	<b>rack</b> 62:20,20 <b>racks</b> 62:19 <b>raise</b> 130:25 134:8 <b>ralligator</b> 87:2,7 87:12,20 88:10,20 88:20	<b>reasonable</b> 38:7 74:15 124:6,7
<b>provision</b> 108:12 109:4 113:13 124:18	<b>questionnaires</b> 121:17 126:4 131:21	<b>ranking</b> 14:12 <b>raw</b> 27:17 <b>reach</b> 115:12	<b>recall</b> 27:20 74:3,5 111:2
<b>provisioned</b> 70:4	<b>questions</b> 6:18 8:9 8:16 9:12,17,21 10:7,9 12:18 27:8 62:6 84:15 86:21 117:19 118:8,21 119:1 120:7,11 121:16 124:5 125:11,12,24 126:19,25 133:8	<b>read</b> 16:6,10 29:22 31:21 35:12 73:9 75:17 81:20 103:17 106:4 120:4 121:4,18 122:6,7,14 123:9 123:16,21 131:3,3 131:4	<b>received</b> 10:18 52:12 100:4
<b>provisions</b> 110:18		<b>readable</b> 12:6,6 103:16,16,20 104:1,5,6,11,19,25 105:8	<b>receiving</b> 51:14
<b>pst</b> 135:5		<b>reader's</b> 72:18	<b>recess</b> 68:12
<b>public</b> 22:8 59:3 116:2		<b>reading</b> 104:3 122:10 123:10	<b>recognize</b> 74:4
<b>publicly</b> 22:3,18 28:12		<b>readstate</b> 90:8	<b>recognizing</b> 118:16
<b>puma</b> 85:19		<b>ready</b> 90:19,25	<b>recollection</b> 76:20
<b>purpose</b> 27:25 56:12 78:15 86:3 112:3 122:12 133:1		<b>real</b> 12:18 16:11 38:17 40:11 96:12	<b>recommendation</b> 50:21 71:24
<b>purposes</b> 33:10 41:1 54:23		<b>reality</b> 35:19 37:4 47:5 61:25 72:4 91:2	<b>recon</b> 91:20
<b>push</b> 56:10		<b>really</b> 16:8 22:6 26:7 27:7 28:1 29:8 45:11 49:4 70:11 72:12 73:6 77:15 83:16,23 97:11 99:13,22 104:9,20 106:17	<b>reconciliation</b> 91:24
<b>put</b> 12:21 26:21 71:17 88:12 95:1 107:6 130:7			<b>reconsideration</b> 134:4,6,17
<b>pymk</b> 86:11			<b>record</b> 6:4 13:1 68:11,13,14 71:23 120:12 123:11 125:4 130:13 133:3,6,9 135:3 136:11
<b>python</b> 59:4,7			<b>records</b> 118:17
<b>q</b>			<b>recovery</b> 49:19
<b>qualification</b> 26:16			<b>red</b> 116:19
<b>qualified</b> 77:16 99:12			<b>redundant</b> 11:6
<b>queries</b> 40:18 41:23 60:7 124:22			<b>refer</b> 43:2 46:1
<b>query</b> 60:7 108:21			<b>reference</b> 17:23 20:6
			<b>referenced</b> 70:2



## [references - rsc]

<b>references</b> 78:22	71:16 136:1,5	<b>returning</b> 39:20	92:8,15,20,20
<b>referred</b> 39:10	<b>reporting</b> 40:4	<b>reveal</b> 130:11	93:10 94:25 95:6
<b>referring</b> 47:24	48:20	<b>review</b> 97:14	95:10 96:14 97:2
50:5 72:4 76:7	<b>reports</b> 36:16	117:2 128:25	97:3,17,21 100:5
77:25 90:12	<b>repository</b> 54:15	129:25 130:1	100:12 101:11,13
115:24 133:21	55:9 114:2	<b>reviewed</b> 126:13	101:20,24 102:20
<b>reflect</b> 103:7	<b>request</b> 20:17	<b>reviewing</b> 10:6	102:21 103:7,13
<b>regarding</b> 6:20	33:19 81:23 84:4	133:16	104:1,2,6,8,9,22
<b>regulatory</b> 15:25	120:13 129:2	<b>revisit</b> 83:13,24	106:9 107:7 110:7
48:20	130:9 134:5	84:1	110:10,20 111:21
<b>reiterate</b> 122:11	<b>requests</b> 44:6	<b>right</b> 7:7,12,21,22	113:9,22,22
<b>relate</b> 50:8 127:22	<b>required</b> 30:18	8:1,25 9:13 10:17	114:22 116:2,8
<b>related</b> 38:3 48:21	130:23 133:7	12:17 13:3,8 14:7	117:1,7 118:12
60:14 74:25 75:3	<b>requirements</b> 9:6	15:14 16:1,3 17:5	119:15 121:25
100:17	9:7 13:23 15:25	17:7,23 20:7,10	122:13,14,23
<b>relates</b> 74:10	109:9	21:23 22:21,22	124:17,20 127:25
<b>relational</b> 9:1	<b>requires</b> 107:1	23:9,20 24:6,9,13	129:5 131:23
108:9 109:2 112:5	<b>research</b> 58:15,16	25:8,14 26:1,3,9	133:14 134:23
<b>relationship</b> 14:17	92:21 93:4,5,8	27:1,15,21 28:2,8	135:2
<b>relationships</b>	<b>researchers</b> 92:23	28:12 32:9,17	<b>rights</b> 123:22
116:15	92:24	34:1,2,4,9 35:23	<b>risk</b> 48:6,19 60:22
<b>relative</b> 136:16	<b>reserve</b> 123:22	36:11,25 37:6,9,10	<b>road</b> 3:18
<b>relevant</b> 104:13,20	<b>resolves</b> 124:7	38:9 39:23 40:17	<b>roadmap</b> 131:5
<b>reliable</b> 92:9 93:16	<b>respect</b> 107:10	40:18,19 41:7,11	<b>rocksdb</b> 23:15
<b>relocate</b> 28:17	<b>respective</b> 66:15	42:1 43:13,19	93:14,19
<b>remember</b> 41:5	<b>respond</b> 33:21	46:8 47:8,19	<b>rohrback</b> 3:10
74:2,8 100:16,18	131:13 133:4	48:12 49:9,13	<b>rolling</b> 116:21
<b>remote</b> 1:13 2:10	134:5	50:1 51:8 54:7,11	133:17
<b>removed</b> 129:6	<b>responding</b> 105:21	55:24 57:11,18	<b>room</b> 10:24 117:6
<b>repeat</b> 94:19	134:16	58:14 60:18,19	117:18,23
<b>repeatedly</b> 119:22	<b>response</b> 17:3,3	61:1,16 62:4,10,17	<b>rooms</b> 66:16
<b>repeating</b> 17:14	29:22,22 87:11	63:23 64:4,7,23	<b>root</b> 83:15
39:7	<b>responsible</b> 88:1,3	66:10,25 68:9	<b>rough</b> 129:9,10
<b>replays</b> 96:12	88:8	69:11 70:2 71:11	131:3
<b>repo</b> 55:5,5	<b>rest</b> 121:23	71:13,21 73:22	<b>route</b> 20:8
<b>report</b> 25:12	<b>restating</b> 113:10	74:21,21 76:14,14	<b>rows</b> 9:3 109:8,21
133:13,14,17	<b>restoring</b> 89:24	76:18 78:2 79:12	<b>royal</b> 30:2
<b>reported</b> 1:18	<b>results</b> 36:9	79:18 80:7 81:1	<b>rpm</b> 116:18
<b>reporter</b> 2:11 6:6	<b>resume</b> 121:1	82:3,17,22,22,25	<b>rpr</b> 1:18 136:25
6:10 39:17,18	<b>retain</b> 7:14	83:19,21,22 89:10	<b>rsc</b> 94:7
40:22,23 70:25		89:20 90:7 91:19	

## [rules - signature]

<b>rules</b> 130:22 <b>ruling</b> 130:1,11 <b>run</b> 25:11 27:7 33:23 44:3 50:2 98:24 110:17 <b>running</b> 33:4 41:23 53:12 58:4 58:11,15 75:17 106:23 108:14,19 112:9 <b>runs</b> 25:13 32:22 43:23 67:16 88:11 99:25 <b>rush</b> 129:10	104:25 107:20 109:20 110:22 116:22 122:15,20 122:22 <b>scale</b> 46:17 <b>scenario</b> 29:18 <b>schedule</b> 25:1,9,12 29:24 <b>scheduler</b> 24:24 <b>scheduling</b> 129:21 <b>scope</b> 13:20 15:12 21:8 126:20,21 127:4 132:17 134:19 <b>score</b> 35:24 <b>screen</b> 7:2,9 <b>scribe</b> 94:10 <b>scripts</b> 36:15 <b>scroll</b> 10:17 18:24 <b>scuba</b> 94:13,14 <b>se</b> 21:19 30:1 118:10 <b>seamless</b> 127:15 <b>seattle</b> 3:12 <b>sec</b> 19:2 <b>second</b> 84:1 121:1 <b>section</b> 13:25 <b>secure</b> 27:3 125:1 <b>securedumbo</b> 95:6 <b>security</b> 114:4 124:20 <b>see</b> 7:9,11,16,17,21 11:13 20:4 46:8 65:16 80:10 127:13 128:11,19 128:25 133:18 <b>seen</b> 90:11 126:10 <b>segments</b> 42:5 <b>select</b> 89:3 <b>selection</b> 33:20	<b>self</b> 108:8 109:1 112:11 <b>selling</b> 99:2 <b>send</b> 9:17 23:22 29:21,21 70:25 80:9 81:18,19 <b>sending</b> 51:12 56:7 71:20 <b>sense</b> 35:8 39:6 58:13 73:1,2 107:7 113:23 127:23 <b>sensitive</b> 42:14 <b>sent</b> 78:19,24 81:25 82:5 92:4 105:23 108:18 126:3 131:17 <b>sentences</b> 104:17 <b>separate</b> 47:3 71:6 117:5 <b>separately</b> 130:25 <b>serve</b> 112:11 <b>server</b> 62:20 68:23 <b>serves</b> 113:17 <b>service</b> 16:11 23:21,24 24:12 45:7 48:14 65:19 72:11,13,21,23 85:8,11 93:15 94:21 96:6,24 99:21,24,25 102:1 106:8 108:8,13 109:1 110:18 <b>servicefriend</b> 95:11 <b>servicelab</b> 95:14 <b>services</b> 24:2 37:5 41:2 48:16 73:10 88:23 112:9 <b>serving</b> 112:7	<b>set</b> 22:19 23:5 31:19,21,22 42:20 109:3 113:12 128:5 130:23 134:25 136:8 <b>sets</b> 35:14 <b>setting</b> 109:19 <b>settings</b> 22:19 28:12 <b>sev</b> 45:25 46:2 <b>seven</b> 77:8,24 <b>sftp</b> 95:18 96:6 <b>shadowdb</b> 96:8 <b>share</b> 7:2,3,9 15:7 36:10,10,13 61:14 67:23 75:3 82:7 110:3 113:5 <b>shared</b> 12:3,5,8 13:10,14,17 15:4 16:8 18:9 21:19 33:12 39:13 49:1 49:1 57:7 62:8 63:3 64:8,22 72:25 75:21 81:15 85:21 89:16,18,19 90:15 99:10 101:6 101:14 111:18 112:13 120:21 122:19,20,25 <b>shares</b> 67:14 97:18 100:8,23 <b>sharing</b> 22:14 63:25 70:20 <b>short</b> 52:22 109:17 <b>shorthand</b> 2:11 6:10 136:4,12 <b>side</b> 85:1 132:14 <b>sigma</b> 100:13 <b>signals</b> 44:16 <b>signature</b> 136:23
<b>s</b>			
<b>s</b> 5:8 <b>s3</b> 67:9,17 <b>saber</b> 35:10 <b>salient</b> 128:23 <b>salted</b> 10:22 <b>sampled</b> 54:22 <b>saw</b> 18:25 <b>saying</b> 12:23,25 33:16 50:12 71:4 96:2 125:14 <b>says</b> 7:12 10:21,22 14:21 16:10,14,15 16:17 18:6,7 19:12 23:5,11,12 26:1 27:9 36:9,18 37:13,14 41:12 42:4,10 45:17 48:5,6 52:3 55:19 59:10 60:6 62:18 63:11,12 65:1,16 67:5 69:25 70:17 72:15 73:17 77:17 83:14 93:14,21 95:19 96:11 99:14 99:16 101:5 102:10 103:2			

**[similar - streams]**

<b>similar</b> 68:1 93:16 98:14 116:20 <b>simple</b> 8:16,23 9:24 90:4 119:6 119:16 124:25 <b>single</b> 82:4 131:21 <b>sip</b> 114:25 <b>sisyphus</b> 96:19,22 <b>site</b> 46:1,3,10 <b>sits</b> 48:17 60:17 73:25 84:25 102:19 113:8 <b>sitting</b> 112:15 119:7 124:9 <b>slash</b> 61:2 <b>small</b> 23:5 69:25 <b>smith</b> 78:9 <b>social</b> 11:4 14:22 28:25 37:8 38:4 101:2 102:11,15 102:17,19,23 111:12 112:17,22 112:25 113:3,7,8 113:20 <b>software</b> 97:25 <b>solely</b> 48:21 <b>solution</b> 13:9 23:7 47:25 88:17 106:16 <b>solving</b> 72:18 <b>somebody</b> 9:16 77:11 <b>something's</b> 13:16 <b>somewhat</b> 108:10 <b>soon</b> 23:7 <b>sorry</b> 11:18 14:23 19:2 20:25 38:18 49:6,9 64:4 79:16 94:18 133:20 <b>sort</b> 8:5 40:10 125:3	<b>sought</b> 117:21 <b>sound</b> 17:14 34:13 <b>sounded</b> 132:15 <b>sounds</b> 7:7 59:14 59:17 84:21 124:4 124:13 127:16 <b>source</b> 54:15 118:20 119:7,7 125:23 <b>sources</b> 80:19 118:22,23 125:10 130:10 132:15,22 134:18 <b>space</b> 36:3 48:2 111:3,20 <b>sparse</b> 97:8 <b>sparta</b> 97:13 <b>speak</b> 49:4 99:22 104:9 117:6 <b>speaks</b> 109:23 <b>special</b> 1:11 22:4 133:10 <b>specific</b> 11:22 16:5 29:2,4 39:2 43:20 45:5 77:19 78:8 80:4,5,20 82:1 83:3 106:18 125:10 132:2 <b>specifically</b> 52:21 83:4,9 123:6 <b>specificity</b> 60:17 <b>specifics</b> 69:21 <b>specify</b> 132:1 <b>speculate</b> 16:21 83:6 <b>speculating</b> 38:7 45:12 46:18 86:9 105:16 <b>spellings</b> 71:9 <b>spent</b> 7:15	<b>spikes</b> 34:1 <b>spin</b> 113:1 119:21 <b>spoke</b> 134:9 <b>spreadsheet</b> 5:11 10:18 17:2 71:8,9 <b>spreadsheets</b> 128:18 <b>srt</b> 63:18 <b>stack</b> 92:18 93:20 <b>stand</b> 45:25 116:18 <b>standalone</b> 37:14 <b>standard</b> 2:13 81:24 <b>stands</b> 111:8 <b>start</b> 10:16 133:14 133:23 134:1 <b>started</b> 6:5 10:5 <b>starting</b> 133:18 <b>state</b> 2:12 113:10 120:12 136:5 <b>stated</b> 46:9 60:18 76:14 119:13,22 <b>statement</b> 108:10 <b>statements</b> 133:8 <b>states</b> 1:1 2:1 <b>static</b> 19:17 <b>status</b> 129:12,15 129:18 <b>stay</b> 117:3,10,25 <b>stays</b> 33:12 47:19 <b>step</b> 89:15,21 123:4 124:9 <b>steroids</b> 95:5 <b>storage</b> 8:15,25 19:19 20:14 21:15 23:15 27:17 30:20 31:7 38:25 39:1 42:15 43:11 47:16 48:12,24 49:21 59:22 60:13 67:20	70:4,13 72:17 73:10 85:8,11 88:3,23 92:12,18 92:19 93:17 94:3 94:7 98:6 101:23 102:2,12,18 106:11,12,13,15 106:16,19,20,23 106:25 107:1 108:22 <b>store</b> 8:17 10:23 10:23 15:16 24:25 26:15,19,22,25,25 38:9,10 39:2 52:10 54:20 63:13 72:17 73:12,25 85:1 86:22,25 88:23,25 94:2 109:2 112:6,11 115:3 127:7 <b>stored</b> 8:18 9:25 10:23 15:18 19:20 21:9 37:16 49:22 52:12 55:11 60:4 63:16 67:12,18,19 78:12 87:21 89:14 99:9 107:3 112:12 113:5 <b>stores</b> 49:25 62:25 <b>stories</b> 90:10,12 <b>storing</b> 70:18 72:6 78:16 93:1 100:2 <b>strategy</b> 13:21 <b>stream</b> 64:16 85:25 86:4 97:23 98:1 <b>streaming</b> 64:4 <b>streamline</b> 9:21 17:16 121:3 <b>streams</b> 16:13
--	---	---	---

**[street - teachers]**

<b>street</b> 3:6 116:2	90:17 92:25 96:17	68:19 69:4,13	<b>tables</b> 40:15,16,18
<b>string</b> 83:20	97:20 98:18,22	70:17 71:7,25	40:19 108:14
<b>strow's</b> 133:13	99:5 101:24	72:9 73:4 75:1,4,5	109:4 124:15
<b>structure</b> 53:20	105:25 106:1	75:16,21 76:8	<b>taco</b> 99:17
<b>studio</b> 90:20 91:2	107:22 116:3	77:6,13 79:4,24	<b>tacosrs</b> 99:4
<b>stylus</b> 97:21,22	121:7,10 125:2	81:25 82:6 85:6	<b>tag</b> 36:22
111:3	126:13 127:14	86:24 87:5 88:1	<b>tailer</b> 86:11
<b>subject</b> 13:13	133:13	88:11,22 91:14,22	<b>tailored</b> 131:22
<b>submission</b> 29:12	<b>surprised</b> 102:6	92:10 93:2 94:3	<b>take</b> 27:2 33:24
<b>submit</b> 6:22 69:6,9	<b>surprises</b> 66:17	94:24 95:13 97:11	66:14 67:6 68:10
<b>submitted</b> 6:19	<b>survey</b> 78:19	98:12,15 99:9	82:23 115:22
17:20 29:11 81:1	81:18,19	106:8,22 111:10	131:1
130:6	<b>surveying</b> 8:13	113:15,16 114:14	<b>takedown</b> 63:18
<b>subscribed</b> 136:19	<b>surveys</b> 78:23	114:19 115:1,2,17	<b>taken</b> 2:10 68:12
<b>subset</b> 55:21 76:12	<b>swift</b> 98:3,10	116:5,7 119:18	136:7
125:6 126:3,11	<b>sworn</b> 6:9	126:6,8 127:19	<b>takeout</b> 24:14
129:2	<b>synapse</b> 98:14	<b>systems</b> 7:13 8:8	<b>takes</b> 19:17
<b>sufficient</b> 126:16	<b>sync</b> 73:10	8:15,19 9:15,22	<b>talk</b> 29:23,23
<b>suite</b> 3:6,12	<b>synchronously</b>	10:12,16 11:22,25	30:11 100:3
<b>supplemental</b>	33:22	12:2 14:6 15:1	117:10,18 120:1
129:14	<b>system</b> 9:24 11:11	18:12,17 21:6	127:23 128:13,14
<b>support</b> 21:16	11:14 12:8 13:5,5	29:20 30:1 39:7	128:16
22:13 73:10 77:4	13:15 14:16,22,22	44:24 50:19 55:18	<b>talked</b> 91:3 107:10
87:16 116:12,13	15:6,7,14,16,19	60:18 72:17 79:22	111:4
<b>supporting</b> 76:16	16:25 17:6,9 18:8	87:16,17,18 88:18	<b>talking</b> 23:14 25:9
<b>supports</b> 109:19	19:18,19,19 20:14	92:12 93:18	81:4 84:23 104:7
<b>sure</b> 6:20,21 7:4	23:2,4,20 24:8	101:22 106:17	104:15 105:18
15:6,10 16:13	26:18 27:2,10,14	116:23,23 119:18	122:16 126:14
23:12,14 26:24	29:5 30:5 31:25	121:16 123:15,21	<b>tally</b> 99:20 100:3
29:13 31:24 34:6	32:2,22 33:9 35:4	124:12,13,23	100:13
35:20 41:18 42:23	35:17 37:2,16	125:6 126:3,15	<b>tao</b> 100:19,24
44:19 46:23 48:10	38:20,25 39:1,2,22	129:2 132:1	102:10,10,19
53:24 58:1 60:9	40:3 42:9 43:10	<b>t</b>	<b>target</b> 26:2
60:10 61:25 62:11	44:22 46:20 47:17	<b>t</b> 5:8	<b>targeting</b> 14:12
63:15,17,19 64:24	48:1,17 49:15	<b>t's</b> 81:3	<b>task</b> 115:5,8,24
66:6,23 74:14,22	50:25 51:2,6,15	<b>tab</b> 67:6	128:7
75:13 79:15,17	52:13,17,19 56:5	<b>table</b> 9:2 109:4,21	<b>tasks</b> 115:16,16
80:10,12,15 81:16	57:4,22 58:3,20	124:17	<b>tax</b> 106:6
81:21 83:3 84:3	60:25 61:8,18	<b>tableau</b> 98:20,21	<b>taxes</b> 106:9
84:23 85:13,17,22	62:2 63:6,16,21,25	98:23 99:1	<b>teachers</b> 49:7
86:2,15 87:5	64:6 65:4,6 67:14		

[team - top]

<b>team</b> 7:12 10:5 11:23 19:12 30:1 30:3 35:6 46:7 61:5 70:3,5 77:4 124:20 131:22	<b>testing</b> 38:2 103:3 103:3 <b>thank</b> 6:16 43:3 117:23 118:3,3 121:12,14,25 122:1,5 123:10 129:20,23 134:25	35:13 37:18 41:4 41:13 42:6 46:22 47:18 50:15 53:4 53:5 56:1 57:7 59:25 60:2,3 62:7 63:4 64:1,8,20 65:2,15,22 67:24 68:7 69:10,17 70:7 72:13,22 73:13 74:20 75:10 75:21 77:11,14 85:8,21 87:3 89:16,19,23 90:5 90:15 91:12,24 92:24 93:11 94:4 95:8 96:20 97:5 97:18 98:21 99:8 100:6 106:13 107:17 110:11 111:18 112:14 113:5 114:13 116:4 118:25 119:4 122:25	<b>ties</b> 101:2 <b>time</b> 2:13 6:17 10:4 16:11 17:15 28:19 30:9 40:11 66:12,21 74:5 75:5 80:14 96:12 109:17,18,25 117:14 118:13 121:21 130:17,19 130:20 136:8 <b>timelinedb</b> 101:13 <b>timing</b> 132:14 <b>today</b> 66:21,22 77:8,24 109:13 120:21 132:16 <b>tokenized</b> 81:17 <b>told</b> 78:20 <b>tool</b> 14:14 16:17 20:23 25:25 31:2 31:14,19,21,22 35:13 36:8,9 39:10,11,12,24 41:22 42:12,14,20 42:24 45:17,20 46:10 48:19 49:14 49:23 52:23 53:18 53:25 54:1,1,3,7 59:8,13 61:21 62:18,23 63:8 66:1 74:18 77:4 94:15,18,21 97:14 98:4,4 101:6 113:1
<b>teams</b> 8:14 46:5 59:10 69:25 70:1 70:5,6 115:13,20	<b>theoretically</b> 15:1 <b>thereof</b> 136:14 <b>thing</b> 26:21 <b>things</b> 8:3,9 10:2 40:6 79:7 80:2 82:13 94:24 112:25 117:20	65:2,15,22 67:24 68:7 69:10,17 70:7 72:13,22 73:13 74:20 75:10 75:21 77:11,14 85:8,21 87:3 89:16,19,23 90:5 90:15 91:12,24 92:24 93:11 94:4 95:8 96:20 97:5 97:18 98:21 99:8 100:6 106:13 107:17 110:11 111:18 112:14 113:5 114:13 116:4 118:25 119:4 122:25	<b>timelinedb</b> 101:13 <b>timing</b> 132:14 <b>today</b> 66:21,22 77:8,24 109:13 120:21 132:16 <b>tokenized</b> 81:17 <b>told</b> 78:20 <b>tool</b> 14:14 16:17 20:23 25:25 31:2 31:14,19,21,22 35:13 36:8,9 39:10,11,12,24 41:22 42:12,14,20 42:24 45:17,20 46:10 48:19 49:14 49:23 52:23 53:18 53:25 54:1,1,3,7 59:8,13 61:21 62:18,23 63:8 66:1 74:18 77:4 94:15,18,21 97:14 98:4,4 101:6 113:1
<b>tech</b> 92:18 <b>technical</b> 88:6 <b>technologies</b> 88:19 <b>technology</b> 43:8 59:3 86:14 102:5 112:11	<b>think</b> 9:20 15:2,2 16:22 17:8 18:25 18:25 29:14 30:8 33:15 45:21 51:7 52:21 53:23 57:15 61:13 64:2 66:16 67:15 74:14,15 76:7 77:11 78:12 80:1 83:12 85:14 88:15 98:11 99:7 99:11 100:2 101:9 103:22 104:13 105:2 106:3 111:8 114:6 116:19 117:8 118:8,15,21 120:10 123:15,17 123:23 125:21 127:9 128:2,24 129:7 130:11,16 131:2 132:8 134:15,25	<b>thought</b> 40:9 42:2 45:7 49:6 62:10 62:17 81:3 114:16 117:4 127:10 <b>thoughts</b> 86:7 <b>thousands</b> 112:8 115:16,17,21 <b>threat</b> 101:5 <b>three</b> 32:2 123:17 128:16 131:6 <b>thrift</b> 67:5 <b>ticketing</b> 115:17 <b>tickets</b> 115:18 <b>tie</b> 83:2 103:1 <b>tied</b> 83:4 102:2 <b>tier</b> 32:23 33:17 42:10 44:2,3	<b>tooling</b> 115:15 <b>tools</b> 35:11 40:5,8 49:12 53:22 <b>top</b> 38:24 39:2 48:17 67:16,20 84:25 99:25 106:17,23 108:14 112:9,15
<b>tel</b> 3:7,13,19,23 20:20 21:21 22:12	<b>thinking</b> 83:9,16 124:24 <b>third</b> 3:12 12:5,8 13:10 18:11 21:20 21:24 24:4 26:6 27:11,12 32:14	<b>thought</b> 40:9 42:2 45:7 49:6 62:10 62:17 81:3 114:16 117:4 127:10 <b>thoughts</b> 86:7 <b>thousands</b> 112:8 115:16,17,21 <b>threat</b> 101:5 <b>three</b> 32:2 123:17 128:16 131:6 <b>thrift</b> 67:5 <b>ticketing</b> 115:17 <b>tickets</b> 115:18 <b>tie</b> 83:2 103:1 <b>tied</b> 83:4 102:2 <b>tier</b> 32:23 33:17 42:10 44:2,3	<b>tooling</b> 115:15 <b>tools</b> 35:11 40:5,8 49:12 53:22 <b>top</b> 38:24 39:2 48:17 67:16,20 84:25 99:25 106:17,23 108:14 112:9,15
<b>tell</b> 8:5 10:13 11:10 16:6 27:24 34:17 43:5 44:23 47:23,23 73:14 76:5 92:6 121:6 126:5	<b>think</b> 9:20 15:2,2 16:22 17:8 18:25 18:25 29:14 30:8 33:15 45:21 51:7 52:21 53:23 57:15 61:13 64:2 66:16 67:15 74:14,15 76:7 77:11 78:12 80:1 83:12 85:14 88:15 98:11 99:7 99:11 100:2 101:9 103:22 104:13 105:2 106:3 111:8 114:6 116:19 117:8 118:8,15,21 120:10 123:15,17 123:23 125:21 127:9 128:2,24 129:7 130:11,16 131:2 132:8 134:15,25	<b>thought</b> 40:9 42:2 45:7 49:6 62:10 62:17 81:3 114:16 117:4 127:10 <b>thoughts</b> 86:7 <b>thousands</b> 112:8 115:16,17,21 <b>threat</b> 101:5 <b>three</b> 32:2 123:17 128:16 131:6 <b>thrift</b> 67:5 <b>ticketing</b> 115:17 <b>tickets</b> 115:18 <b>tie</b> 83:2 103:1 <b>tied</b> 83:4 102:2 <b>tier</b> 32:23 33:17 42:10 44:2,3	<b>tooling</b> 115:15 <b>tools</b> 35:11 40:5,8 49:12 53:22 <b>top</b> 38:24 39:2 48:17 67:16,20 84:25 99:25 106:17,23 108:14 112:9,15
<b>temporarily</b> 63:16 <b>temporary</b> 51:13 76:25	<b>thinking</b> 83:9,16 124:24 <b>third</b> 3:12 12:5,8 13:10 18:11 21:20 21:24 24:4 26:6 27:11,12 32:14	<b>thought</b> 40:9 42:2 45:7 49:6 62:10 62:17 81:3 114:16 117:4 127:10 <b>thoughts</b> 86:7 <b>thousands</b> 112:8 115:16,17,21 <b>threat</b> 101:5 <b>three</b> 32:2 123:17 128:16 131:6 <b>thrift</b> 67:5 <b>ticketing</b> 115:17 <b>tickets</b> 115:18 <b>tie</b> 83:2 103:1 <b>tied</b> 83:4 102:2 <b>tier</b> 32:23 33:17 42:10 44:2,3	<b>tooling</b> 115:15 <b>tools</b> 35:11 40:5,8 49:12 53:22 <b>top</b> 38:24 39:2 48:17 67:16,20 84:25 99:25 106:17,23 108:14 112:9,15
<b>terabyte</b> 40:15,15 <b>terabytes</b> 40:18	<b>thinking</b> 83:9,16 124:24 <b>third</b> 3:12 12:5,8 13:10 18:11 21:20 21:24 24:4 26:6 27:11,12 32:14	<b>thought</b> 40:9 42:2 45:7 49:6 62:10 62:17 81:3 114:16 117:4 127:10 <b>thoughts</b> 86:7 <b>thousands</b> 112:8 115:16,17,21 <b>threat</b> 101:5 <b>three</b> 32:2 123:17 128:16 131:6 <b>thrift</b> 67:5 <b>ticketing</b> 115:17 <b>tickets</b> 115:18 <b>tie</b> 83:2 103:1 <b>tied</b> 83:4 102:2 <b>tier</b> 32:23 33:17 42:10 44:2,3	<b>tooling</b> 115:15 <b>tools</b> 35:11 40:5,8 49:12 53:22 <b>top</b> 38:24 39:2 48:17 67:16,20 84:25 99:25 106:17,23 108:14 112:9,15
<b>term</b> 35:7 52:22 54:25 55:2 77:21 97:12 98:19 103:16	<b>thinking</b> 83:9,16 124:24 <b>third</b> 3:12 12:5,8 13:10 18:11 21:20 21:24 24:4 26:6 27:11,12 32:14	<b>thought</b> 40:9 42:2 45:7 49:6 62:10 62:17 81:3 114:16 117:4 127:10 <b>thoughts</b> 86:7 <b>thousands</b> 112:8 115:16,17,21 <b>threat</b> 101:5 <b>three</b> 32:2 123:17 128:16 131:6 <b>thrift</b> 67:5 <b>ticketing</b> 115:17 <b>tickets</b> 115:18 <b>tie</b> 83:2 103:1 <b>tied</b> 83:4 102:2 <b>tier</b> 32:23 33:17 42:10 44:2,3	<b>tooling</b> 115:15 <b>tools</b> 35:11 40:5,8 49:12 53:22 <b>top</b> 38:24 39:2 48:17 67:16,20 84:25 99:25 106:17,23 108:14 112:9,15
<b>terms</b> 53:25 61:3 71:10 72:2 84:2,4 93:4 102:14 118:15	<b>thinking</b> 83:9,16 124:24 <b>third</b> 3:12 12:5,8 13:10 18:11 21:20 21:24 24:4 26:6 27:11,12 32:14	<b>thought</b> 40:9 42:2 45:7 49:6 62:10 62:17 81:3 114:16 117:4 127:10 <b>thoughts</b> 86:7 <b>thousands</b> 112:8 115:16,17,21 <b>threat</b> 101:5 <b>three</b> 32:2 123:17 128:16 131:6 <b>thrift</b> 67:5 <b>ticketing</b> 115:17 <b>tickets</b> 115:18 <b>tie</b> 83:2 103:1 <b>tied</b> 83:4 102:2 <b>tier</b> 32:23 33:17 42:10 44:2,3	<b>tooling</b> 115:15 <b>tools</b> 35:11 40:5,8 49:12 53:22 <b>top</b> 38:24 39:2 48:17 67:16,20 84:25 99:25 106:17,23 108:14 112:9,15
<b>test</b> 54:23 69:6,6 96:9,12	<b>thinking</b> 83:9,16 124:24 <b>third</b> 3:12 12:5,8 13:10 18:11 21:20 21:24 24:4 26:6 27:11,12 32:14	<b>thought</b> 40:9 42:2 45:7 49:6 62:10 62:17 81:3 114:16 117:4 127:10 <b>thoughts</b> 86:7 <b>thousands</b> 112:8 115:16,17,21 <b>threat</b> 101:5 <b>three</b> 32:2 123:17 128:16 131:6 <b>thrift</b> 67:5 <b>ticketing</b> 115:17 <b>tickets</b> 115:18 <b>tie</b> 83:2 103:1 <b>tied</b> 83:4 102:2 <b>tier</b> 32:23 33:17 42:10 44:2,3	<b>tooling</b> 115:15 <b>tools</b> 35:11 40:5,8 49:12 53:22 <b>top</b> 38:24 39:2 48:17 67:16,20 84:25 99:25 106:17,23 108:14 112:9,15
<b>testified</b> 6:11 <b>testifying</b> 136:10 <b>testimony</b> 132:15	<b>thinking</b> 83:9,16 124:24 <b>third</b> 3:12 12:5,8 13:10 18:11 21:20 21:24 24:4 26:6 27:11,12 32:14	<b>thought</b> 40:9 42:2 45:7 49:6 62:10 62:17 81:3 114:16 117:4 127:10 <b>thoughts</b> 86:7 <b>thousands</b> 112:8 115:16,17,21 <b>threat</b> 101:5 <b>three</b> 32:2 123:17 128:16 131:6 <b>thrift</b> 67:5 <b>ticketing</b> 115:17 <b>tickets</b> 115:18 <b>tie</b> 83:2 103:1 <b>tied</b> 83:4 102:2 <b>tier</b> 32:23 33:17 42:10 44:2,3	<b>tooling</b> 115:15 <b>tools</b> 35:11 40:5,8 49:12 53:22 <b>top</b> 38:24 39:2 48:17 67:16,20 84:25 99:25 106:17,23 108:14 112:9,15



[touchy - user]

<b>touchy</b> 13:12	<b>turns</b> 132:10	125:14,21 126:14	93:7,24,25 96:3
<b>trace</b> 11:13	<b>tw</b> 101:23	128:6 134:22	98:7 99:12,19,22
<b>track</b> 59:10 114:4	<b>two</b> 12:18 34:24	<b>understandable</b>	100:2 113:15
115:18 125:1	37:22 100:17	104:19	126:13 132:5
<b>tracking</b> 61:3	117:6 122:8	<b>understanding</b>	<b>user</b> 1:6 2:6 7:14
124:21	128:16 134:12	18:21 36:2 62:1	8:17,18,22,22,24
<b>traffic</b> 19:13 96:11	<b>type</b> 67:17 78:12	62:21 76:17 90:3	9:2,3,7,25 10:13
96:16	81:11 93:8 104:18	103:4 104:11	11:3,12 12:2,19,20
<b>training</b> 16:16	127:2,8	114:23 117:13	14:1 15:5,16,18
114:12	<b>typical</b> 20:3	123:23 127:1,21	18:9 19:15 20:1
<b>transaction</b> 70:12	<b>u</b>	128:8 131:19,25	20:17,18 21:9,16
74:23	<b>udb</b> 102:8,11	<b>understood</b> 80:16	21:19 22:17,25
<b>transactional</b> 75:9	<b>uh</b> 7:20 10:20	80:17 98:3	24:20,25 25:1
75:11	22:16 47:14 48:25	<b>underway</b> 133:25	26:4,7,13,15,19,25
<b>transcribed</b>	55:16 62:5 69:16	<b>unduly</b> 114:17	26:25 27:1 28:3,5
136:13	72:20 73:20,20	<b>unfortunately</b>	32:7,11,12,24 33:7
<b>transcript</b> 40:21	78:21 79:21 90:21	17:11 111:14	33:9,10,19 37:16
117:2 120:5,13	92:22 100:20	<b>unicorn</b> 14:18	37:17 39:12,25,25
121:4,18 122:6,15	107:9 108:4	<b>uniform</b> 82:4	43:9,12 44:21,22
123:16 128:25	<b>ui</b> 41:22	<b>unit</b> 8:25 69:6	44:23,24 45:10
129:9,25 131:3,4	<b>uid</b> 83:20	<b>united</b> 1:1 2:1	48:10,16 49:24,25
<b>transcription</b>	<b>uip</b> 103:13,14	<b>unittestdb</b> 103:3,5	50:8,10,13,15 52:6
136:14	<b>ultimately</b> 13:21	<b>universal</b> 80:7	52:18 53:20 54:17
<b>transform</b> 23:8	67:18 129:1	<b>unstructured</b> 65:9	54:21 55:10 56:19
<b>translate</b> 104:2	<b>unauthorized</b>	<b>up2x</b> 103:10,12	57:6,14,24 58:18
<b>translates</b> 103:25	125:2	<b>update</b> 49:8	60:17,23 63:2,15
<b>tru</b> 101:19	<b>underlying</b> 19:18	<b>upload</b> 20:18	63:16 64:13,13
<b>true</b> 105:5 109:6	43:11 102:15	22:18	65:6,13 66:6
119:17 126:7	106:19	<b>uploading</b> 79:7	67:11,14 68:24,25
127:11	<b>underneath</b> 93:19	<b>uploads</b> 22:17	69:7 70:18 72:16
<b>truly</b> 103:6	<b>undersigned</b> 136:4	<b>upset</b> 17:17	74:9,10 75:11,22
<b>try</b> 41:21	<b>understand</b> 7:14	<b>usage</b> 77:24	76:5,6,10,10,18,23
<b>trying</b> 10:12 17:19	8:4,21 9:20 10:12	<b>use</b> 11:23 13:13	77:13 78:5,6,6,7
18:24 26:2 74:2	14:21 21:3 29:16	20:3 22:10 29:8	78:20,25 79:1,2,10
80:18,23 104:10	29:25 37:8 53:1	29:10,11,25 30:4,7	79:11,18,23,25
105:19 124:25	57:20 58:14 63:2	30:11,14,19 33:8	80:13,19 81:2,4,11
125:25 127:14	65:10 67:11 68:25	35:14 38:23 46:5	81:21 82:8,10,11
128:12,20 131:7	79:8 95:3 101:21	49:4 51:6 52:8	82:14,16,20 83:4
<b>tupperware</b> 102:1	104:1 105:19	55:7 60:14 66:14	83:11,16,22 85:1
<b>turn</b> 118:1	113:11 117:21	78:13 79:8,9 80:2	85:15,16 86:15,17
	118:13 120:1,20	83:19 86:18 87:18	86:25 89:1,8,13,16

[user - worked]

90:9,10,22 91:8,11 91:16 92:5 93:1,5 95:25 96:1,14,15 96:15,15 98:23,25 99:9 100:24 104:24 105:9,10 105:18,20,24 106:12 107:1,3 108:17,18 109:6,9 109:13 111:21 112:12 115:6,23 116:4 119:17 124:23 125:3 126:6,6,7 127:2,5 127:7,18,22,24 128:12 132:5,6 <b>user's</b> 28:24 72:24 81:13 108:23 <b>users</b> 10:22 21:25 24:17 35:24 54:16 56:7 72:7 76:12 76:15 77:22 78:8 79:6 82:10,12 83:17 91:4,25 105:25 123:3,6 <b>uses</b> 13:6 23:16 51:25 53:19 60:15 93:19 123:3 <b>ux</b> 92:23	<b>version</b> 68:23 69:4 76:8 111:3 <b>versus</b> 82:1 104:17 <b>vett</b> 82:24 <b>viable</b> 35:2 <b>video</b> 37:15 64:7 64:12,13,14,15,18 91:16 118:1 <b>videoconference</b> 1:14 3:9,14,20,24 4:4,7 <b>videos</b> 19:20 22:14 49:22 64:17 <b>view</b> 15:13 70:3 <b>virtual</b> 35:19,19 37:4 45:4,4 <b>visibility</b> 78:17 <b>vision</b> 114:18,21 <b>visual</b> 94:15 <b>volumes</b> 92:9 93:16 <b>vr</b> 36:18 37:1,3	134:8 <b>wanted</b> 17:20 79:17 109:2 121:12 <b>wants</b> 19:24 20:20 117:9 121:6 <b>warehouse</b> 40:6 59:21 60:12 68:17 110:13 <b>warehouses</b> 124:1 124:14 <b>warehousing</b> 98:9 <b>warm</b> 106:11,12 106:13,15,19,23 106:25 107:1 108:22 <b>washington</b> 3:12 <b>wasting</b> 80:13 <b>wave</b> 115:5,8,8,9 115:22 <b>waves</b> 115:14 <b>way</b> 4:3 8:8,12 19:16 33:4 35:20 40:5 51:4 73:9,18 83:8 86:10 106:21 106:25 107:6,11 108:11,25 110:6 116:20 118:5 124:21,25 126:2 126:20 131:12 <b>we've</b> 61:3 66:11 <b>weaver</b> 3:4 118:12 119:11 120:20 121:10 122:11,22 123:12 129:22,23 131:1 132:9,11,12 132:13 134:15,23 <b>web</b> 32:23 33:8,17 37:2 44:2,3 67:7 111:12	<b>website</b> 37:7 <b>week</b> 133:23 134:2 <b>weekend</b> 121:14 121:23 122:2 <b>weigh</b> 130:4,10 131:11 134:17 <b>welcome</b> 122:1 <b>went</b> 11:23 57:13 75:24 104:14 127:17 <b>west</b> 28:19 66:12 <b>whatsapp</b> 43:9,10 43:20 58:11 107:8 107:10,12,16 <b>wheels</b> 133:25 <b>whereof</b> 136:18 <b>wi</b> 111:9,13 <b>willing</b> 118:9,10 118:11 <b>witness</b> 5:2 6:9 81:22 117:15,22 121:25 122:4 136:18 <b>witnesses</b> 136:9 <b>wondering</b> 39:5 <b>word</b> 54:24 <b>words</b> 37:21 103:18 <b>work</b> 9:19 15:15 16:24 18:2 40:24 45:1 52:6 66:19 66:23 101:9 103:5 104:20 105:4 115:10 120:17 125:23 126:1 128:3,7,21 <b>worked</b> 8:21 9:4 10:4 11:24 12:14 14:18 53:8 73:21 74:5 98:20
<b>v</b>	<b>w</b>		
<b>value</b> 39:1 88:23 112:6,11 <b>variance</b> 78:13 <b>various</b> 8:14 105:9 125:20 <b>vary</b> 88:17 <b>vc</b> 1:6 2:6 <b>vendors</b> 95:20,22 <b>verbatim</b> 136:11 <b>verification</b> 115:23	<b>w</b> 3:11 <b>wait</b> 130:19,19 <b>wallet</b> 48:14,18 <b>want</b> 8:2 9:18 12:7 16:20,23 17:16 18:18 24:10 28:6 32:18 34:5 45:8 50:22 53:24 57:25 66:13 79:15 81:15 83:5,5 84:5,18,22 107:22 109:3 110:17 112:4,5 115:11,12 117:10 117:18 120:2,18 124:5 125:8 126:22 127:13 128:3,5,5 130:21 131:2,11 133:4,12		

## [workflow - zoom]

<b>workflow</b> 41:2,9	<b>y</b>	75:19,23 76:3,21	<b>york</b> 3:23,23
<b>workforce</b> 82:2		77:1,9,19 78:10	<b>z</b>
<b>working</b> 7:15 8:6	<b>y</b> 58:25	79:9,14 80:25	<b>zero</b> 34:20
8:13 30:3 76:21	<b>yeah</b> 7:5 8:11,11	81:22 83:2,5,13	<b>zippy</b> 38:10,24,25
84:8 133:22	9:11,23 11:17	84:2,2,21 85:5,10	39:3 55:24 67:17
<b>workload</b> 34:4,9	12:9 13:19 14:15	85:13,22 86:1,8,13	73:25 84:25 88:23
<b>workloads</b> 98:9	14:18,23 15:10,23	86:19 87:5,13,24	88:25 93:19 99:25
<b>works</b> 19:16 88:14	16:20 17:7,18,25	88:15,15 89:7,7,12	111:15,17,19
108:25 114:19	18:5,13,18,23	90:17,24 91:9,13	112:2,6,9,12 113:2
<b>world</b> 6:16 19:21	19:10,11,22,25	91:15,21 92:1,11	113:15,18,21
28:15 61:24 62:2	20:2,25 21:11,11	92:13,20,25 93:3	<b>zippydb</b> 113:13
<b>wormhole</b> 107:20	21:18 22:3,24	94:8,15,17,20 95:5	<b>zoom</b> 1:14 3:9,14
<b>wrapped</b> 22:19	23:17,19 24:2,15	95:10,12,21 96:9	3:20,24 4:4,7
<b>write</b> 36:15 51:24	24:22 25:16,23	96:17,22 97:10,11	
94:1	26:10 27:18 28:7	97:15,20,22 98:13	
<b>writing</b> 59:4 72:16	29:13 31:4,11,13	98:15,25 99:2,3,11	
<b>written</b> 14:20	31:15 32:20 33:3	99:21 100:10,14	
31:16 46:19 60:7	33:15 34:5,5,5,8	100:16 101:1,8,12	
63:10 64:10,19	34:10,15 35:9	101:18,20,20,25	
68:22 70:14 72:12	36:6 37:3,6 38:6,6	102:4,5,6,9 103:2	
73:7 75:2 88:9	38:12,18,22 39:15	103:9,11,14,19	
91:22 94:9 103:22	39:23 40:3,20	104:14,23 105:2	
<b>wrong</b> 84:20 104:3	41:5,10,20,22	105:12,16,22	
123:20 124:24	42:18,19,24,25	106:5,10,15,22	
<b>wrote</b> 104:10	43:5,12,21,25	107:5,6,9,14,24	
105:17,24	44:17 45:11,14,21	108:2,25 110:6,6	
<b>www</b> 44:1	46:13,18,23 47:1	110:16,21,25	
<b>x</b>	47:22,22 48:23	111:7,14,16,24	
<b>x</b> 5:1,8 111:13	49:4 50:5,17,25	112:2,21,22	
<b>xb</b> 119:18	52:4,4,20 53:12,21	113:19,23 114:6	
<b>xdb</b> 108:3,13,15	54:19 55:2,12,14	114:10 115:1,7,7	
108:19,22,25	55:25 56:16,18	116:6,11 119:11	
109:13,19 110:3	57:3,9,12,12,15,19	120:16 125:25	
110:12,15 112:4,5	58:2,17,20 59:6,21	126:24 128:9	
115:11	60:1 61:20 62:15	133:5 134:21,21	
<b>xstream</b> 110:22,24	62:21 63:9,17,19	<b>year</b> 7:15 8:1	
<b>xt</b> 13:23	64:15 65:16 66:8	34:24 87:9 111:1	
<b>xwf</b> 111:6,8	66:23 67:15 68:5	132:17,19	
	68:21,21 69:2	<b>years</b> 34:24	
	70:2,10,14 73:2,3	<b>yesterday</b> 134:10	
	73:3,15,15,20 74:2		
	74:3,12,14 75:13		



# Exhibit J

GIBSON DUNN

Gibson, Dunn & Crutcher LLP

200 Park Avenue  
New York, NY 10166-0193  
Tel 212.351.4000  
www.gibsondunn.com

Laura C. Mumm  
Direct: +1 212.351.2404  
Fax: +1 212.817.9504  
LMumm@gibsondunn.com

January 6, 2022

VIA JAMS ACCESS

Re: *In re Facebook, Inc. Consumer Privacy User Profile Litigation*, 3:18-md-02843

Dear Special Master Garrie:

Pursuant to the Special Master's Amended Order re: Plaintiffs' Motion to Compel Production of Plaintiff Data, Facebook has enclosed a document containing descriptions of the 149 data systems identified in the Declaration of David Pope as storing or interacting with user data. Consistent with the Special Master's Orders and statements on the record, the descriptions Facebook is providing consist of existing materials prepared by Mr. Pope's team in the course of its work, which was unrelated to this or any litigation.

Facebook has not identified additional existing, readily accessible materials identifying (1) "a description of the types of Named Plaintiff data contained" in each data system; or (2) the "most common functions and purpose" of each data system or "the business units, divisions, or groups" that use the systems. Where Facebook has not been able to provide additional data, Facebook has provided an explanation as to "why this information is not readily accessible," as required by the Order.

Facebook is providing this information pursuant and subject to the terms of the Protective Order that was negotiated by the parties and entered by the Court in this action and has marked this submission Highly Confidential – Attorneys' Eyes Only. These confidential materials are being submitted pursuant to the Special Master's Order for use only in connection with the parties' dispute as to the production of data relating to the Named Plaintiffs, and without prejudice to Facebook's rights and privileges.

The Special Master has also ordered Facebook to make Mr. Pope available ex parte on or before January 14, 2022. Mr. Pope is available from noon PT to 1 pm PT on January 14, 2022.

Respectfully submitted,

/s/ Laura C. Mumm

Laura C. Mumm

GIBSON DUNN

Page **2** of **2**

cc: Lesley E. Weaver

Derek W. Loeser

Enclosure

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
ABP (AWS)	Files with salted and hashed Facebook users' PII are stored in this data store. The data store is used by [REDACTED]	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
AdFox	This is used in the delivery flow for targeting, ranking features, and as information for front-end to display. The social system is not, however, a general purpose system like Unicorn -- it is only for ads. Thus, all objects returned are being actively targeted by an ad.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
AdLogger	AdLogger is a real-time service that performs a join between different event streams and features. AdLogger is the core of the training data preparation flow. It is a realtime joiner of ads events to prepare training data for ads ranking models. The ads events include features, impressions, clicks, conversions, and (mainly) negative-feedback events.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
AdmarketDB (MySQL)	It is the backbone for the Core Ads system not only just Ads buying but also other critical consumers of the Ads system Ads Delivery, Ads Measurement (Reporting), Business integrity, Payments etc.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Ads Raw Storage	Ads Raw Storage is a set of components that together form the platform storage layer in Ads delivery system to provide distributed file system storage. As a common platform layer, this is used by multiple teams across Ads org for various business applications for data storage needs.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Akkio	Akkio Locality is used to improve latencies and reduce replicas for usecases across different datastores. We do this by tracking which regions each ushard is being accessed from and limit the replicas for each ushard to three (or some configurable) number of regions.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
AOS (Ads Online Storage)	Serves ad-side and user-side features for AdFinder, AdPublisher, Prospector, FBShops	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Archival	<div></div> <div>The internal metadata is in MySQL and follows general MySQL backup flow.</div>	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Artillery	Artillery is Facebook's end-to-end distributed tracing system that provides observability into the execution of requests. We support traces across systems (web, mobile, and services) and provide a general purpose platform for near-real time analysis.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Assistant User Memory (AUM)	Assistant provides a <div></div> <div></div> <div>These can be facts implicitly learned from other places in Facebook - user profile, Messenger chats, etc, or explicitly provided by the user to the Assistant as part of a dialog.</div>	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Async	Async is not really a data store per se, it is a platform that allows users to schedule jobs, mostly from code, to be executed outside of the path of a request, thus asynchronously.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Bars	Storing application data for the Bars app (which is a zero to one NPE app that provides a social rapping experience)	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Beat Saber, Multiplay and Analytics	To identify the user, provide scores/leaderboard, connect players to other players, and provide analytics data.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
BeeST and local_namespace	BeeST is a snapshot and shadow testing framework used for functional, correctness and performance testing of Spark, Hive and Presto.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Bento Next	It is used to query and analyze internal data, produce visuals of analyses and share them with others.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.



Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
BigBox VR (M&A)	A system that provides game player and play data support.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Callisto	Callisto is the service which manages the persistent storage of messages for Facebook Messenger. It is the successor of Prometheus and was created in order to migrate from Prometheus' reliance on HBase as the backing store to MySQL with the RocksDB storage engine.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Calypso	Calypso is the project to use the Presto engine as the backend for Atlas Reporting. They are building a Vertica-based solution to deliver to small set of initial customers. The plan is to build out the Presto based solution as soon as possible, and transition onto this platform.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
CDN	CDN is a team within Traffic dedicated to the performance and delivery of media content of all Facebook products such as Facebook, Instagram, WhatsApp, Oculus and more.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.



Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Chronos	Chronos is a distributed scheduler for cron-like jobs (aka bash commands)	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Cocoon	Cocoon is an inventory estimation used to estimate the number of users that can be reached given one or more targeting criteria. It is mainly used by advertisers trying to target their ads to the right audience.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Coefficient	Serves Facebook use cases for user - entity features and scores	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Collab	Powers the Collab application. Collab is a standalone app for creating music video content collaboratively. The data stored in our system is mostly user data (e.g. posts, comments, etc).	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Configurator	Provide tools and storage to author, test, and distribute configuration data	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Connectivity Testing Framework	Provide automated testing of FB developed hardware	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Contributor Store (Zippy)	Contributor Store (Zippy) is used to store contributors and contains data keyed off of the Measurement ID MID) that allows for real-time attribution.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Cortex	Cortex is a digital asset management (DAM) tool, sometimes just referred to as a DAM. Properly speaking a DAM is a digital tool for asset management	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Cosco	Cosco provides efficient disk-based sort-merge for warehouse queries at FB scale.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Cubrick	Cubrick is an in-memory realtime analytics database that enables users to execute sub-second moderately complex SQL queries on tables up to 10+ TB.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
CWS	Core Workflow Services (CWS) is a fault tolerant scalable workflow solution, designed to let people specify a series of steps to be followed and then ensuring the steps lead to an expected conclusion.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Daiquery	Daiquery is an interactive tool for querying multiple data sources at Facebook	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
DAPL	DAPL is the backend platform that Ads Products use to build "Audience Segments".	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Dark Isilon	tier 0 backups needed to recover critical infra & network	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Dataswarm (Gluster FS)	Dataswarm is a data pipeline platform. A data pipeline is a set of data processing operations, where the output of one operation is the input of the next one.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Deld	WhatsApp's Deletion Service	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Deletion Framework	The Deletion Framework is an abstraction in our www codebase for reliably deleting data graphs from online data stores and being able to recover efficiently from errant deletions if they're discovered quickly enough. It's our main system used to satisfy our User Data Deletion Policy for data stores which support point deletes.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Delivery Insights	Generate signals from ad delivery that can be used to provide advertiser guidance	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Dependency Manager Service	It is a Core Systems infra service which allows other infra/product teams build dependencies between FOQS queue items.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
DevDB	DevDB is a MySQL database that runs on a developers devserver.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.



Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
DevOps Investigation Tool (DIT)	The DevOps Investigations team tackles various data-abuse related escalations and instances, ranging from data breaches involving FB data to SEVs enabling third-party apps to access more than they should, to malicious actors monetizing data obtained from the FB platform	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
DIGraph	DIGraph (Data Infrastructure Graph) is a compute engine for large-scale graph analytics.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Downpour Interactive (M&A)	User accounts game player data for Onward from Downpour	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Dumbo	Dumbo is a simple, reliable, low dependency object storage system, with only four high-level operations: put, get, delete, and list.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
dx / Models Framework (Novi)	dx/Models Framework (Novi) is a fast scalable eventually consistent point-in-time structured storage for Novi / F2 risk and compliance assessments	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
EE Storage	The Enterprise Engineering Storage Team is focused on delivering scalable storage solutions to support the applications and virtualization environment that run the internal business at Facebook.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
ElasticSearch	Index database focused on low latency search. Primarily an open source project being supported at Facebook.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Everstore	Everstore provides durable, DR compliant, and highly available online storage. We aim to provide solutions to some of the hardest distributed storage problems in the industry at scale.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	This is one of the various databases that powers the Social Graph. Facebook produced each Named Plaintiff's DYI file, which for each individual user represents the most complete and best compilation of data Facebook maintains associated with that user, and the best available compilation of the data about that user in the Social Graph, in a human readable form.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
F3	The Facebook Feature Framework (F3) is Facebook's next generation feature engineering platform.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
fastpathdb (MySQL)	Iris (aka Generic Iris) is a persistent queue service that guarantees in-order delivery to either devices or backend services.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Fbcode	Fbcode is a top-level directory of the fbsource repo where our backend services code lives. It is a federation of projects, written in a number of languages.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
FBLearner / AI Infra (ML Deletion)	Enables all Facebook engineers to build AI applications and perform AI researches by providing scalable, reliable, efficient, easy-to-use and compliant AI services that power all parts of AI workflow, from feature engineering, to training, to inferencing, during both experimentation and model productionization.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.



Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
FBNS Notification Store (Zippy)	FBNS is a system built to deliver push notifications to mobile apps over MQTT as transport. The overall long term goal is to remove reliance of GCM (Google Cloud Messaging) for pushing notifications.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
fbpkg	Fbpkg is Facebook's scalable binary distribution system.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
FOQS	FOQS powers the Async Tier for job processing. It provides a number of features important to that use case including prioritization of items, a lease expiry for items in the queue, and explicit timing for when an item in the queue is ready for processing.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Forecast	Keeping the NPE app Forecast running ( <a href="https://apps.apple.com/us/app/id1509378877">https://apps.apple.com/us/app/id1509378877</a> )	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
ForgETS	ForgETS is an eventually consistent key-value data store fully implemented in Erlang and designed to run within WA infrastructure. It can run embedded within the backend itself or as a remote tier managed by ForgETS team.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
FORT - FRL Research (AWS)	FORT provides tooling to support external research initiatives. FORT tooling primarily consists of a customized JupyterHub deployment that provides restricted access to FORT datasets.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Goals Tool	For teams to track goals information (objectives, outcomes, metrics for how to measure success)	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Hive	"Hive" refers to several separate systems at Facebook: Hive Tables: Also called Data Warehouse, Hive Query Engine (also called Corona): An engine that executes queries written in FB-HiveQL, and Hive Query Language (HiveQL): A query language data analysis.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	As Facebook has disclosed previously, Hive stores millions of data sets that it uses for internal analytics, product development, and other business functions. As discussed in detail in the Declaration of Mengge Ji, these data sources are not intended to be analyzed or used at the individual user-level and are not structured to allow centralized searches for an individual user's data. Rather, searches for data related to an individual user would be an ad hoc, time-consuming (years long), and complex process that would yield raw data that may not be human readable.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Hive/Risk	Hive/Risk is a data store used to store offline data for analysis	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Hotline / Fireside	Storing user data to deliver a live streaming Q&A type of experience for a stand alone product.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
HulkDB	A general purpose key-value storage solution that can handle high write traffic and store huge data in a highly efficient way.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
IRC	IRC is a resilient plain-text chat protocol widely used by our engineering teams.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Lacis	Lacis is an end-to-end solution for location-anchored content discovery for World AR. At the core is Geo Indexing, a component for retrieving content based on lat/long positions and 2D polygons.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Laser	Laser is a service that indexes data in Hive and Scribe to provide high throughput, low latency lookup. Use this for indexing warehouse data.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
lbu (MySQL)	LBU is a hostname-scheme that exists to run Logtailer Servers. LBUs are designed to be spread thinly across racks to allow that Logtailers can always be allocated in the different racks as the MySQL instances they are tailing from.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
lhubiz (MySQL)	LBU is a hostname-scheme that exists to run Logtailer Servers. LHubiz is a system within this framework. LBUs are designed to be spread thinly across racks to allow that Logtailers can always be allocated in the different racks as the MySQL instances they are tailing from.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Leaf4	Leaf is a fast in-memory storage of all our action/object ranking data. Leaf4 stores action metadata, keyed on actor ID.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Leaf5	Leaf5 is an in-memory cache for actions and objects. More accurately, it is a meta-data cache since it doesn't actually store any physical objects such as photos or posts.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Legal SRT Content Takedown	Reviewing Content that Governments think should get taken down form FB or IG	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
LIV3	Live streaming Infra.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.



Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Live Video Delivery	Deliver live video to viewers	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Live Video Processing	Transcode live video with high quality efficiently, reliably, and low latency.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Lldbpostmortem Cache	Provide a cache for resources available internally at facebook	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Logarithm (previously LogKnock)	LogImport + Logarithm is a new experience to enable reliable realtime storage and efficient search and retrieval for unstructured logging data generated by all applications across Facebook.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
LogDevice	LogDevice is a distributed data store for sequential data. It offers high durability and availability under a variety of workloads and failure scenarios.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Managed Compression Dictionaries	Managed Compression abstracts away the complexities of using dictionary-based compression. It does all the required state management internally (sampling, training, benchmarking, distributing, versioning, garbage collecting, etc.), exposing the same, simple, stateless interface you expect to see from a compression algorithm.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Manifold	Manifold is a managed horizontally scalable BLOB storage service accessible via Thrift and HTTP. It is analogous to Amazon S3.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	This is one of the various databases that powers the Social Graph. Facebook produced each Named Plaintiff's DYI file, which for each individual user represents the most complete and best compilation of data Facebook maintains associated with that user, and the best available compilation of the data about that user in the Social Graph, in a human readable form.
Memcache	Memcache refers to our implementation of memcached, along with its associated client libraries and APIs. Memcache was initially deployed as a distributed in-memory caching layer between the web tier and MySQL UDBs.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Metastore	Managing metadata for Data Warehouse	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Mononoke	Mononoke is the next-generation mercurial-compatible version control server (to scale our monorepo setup)	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Multi Party Live	A thrift service that composites audio and video streams and connects RTC infra with Live Infra.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
MySQL	The MySQL team works to keep all MySQL databases up and running 24x7. We also handle the lifecycle of database assets, and various services to support MySQL at scale.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	This is one of the various databases that powers the Social Graph. Facebook produced each Named Plaintiff's DYI file, which for each individual user represents the most complete and best compilation of data Facebook maintains associated with that user, and the best available compilation of the data about that user in the Social Graph, in a human readable form.



Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
New Product Experiments (AWS)	Small entrepreneurial teams building standalone app experiments.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Novi Gateway Transaction Storage	Novi Gateway Transaction Storage is a queueing system storing user data.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Oculus Recommendation Events	Logging for understanding user interactions on Explore, Store, and other Oculus surfaces	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
OFFD	Storing offline messages	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
OIL	OIL allows you to focus on the problem of "storing/reading a few bytes" at a basic level. It provides the user writing the data with the flexibility to store the bytes on the storage systems of their choice, while solving the reader's "where do I get that data" problem.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Omnistore	Omnistore is a data storage and sync service, which supports periodic polling as well as real time updates to the Facebook family of apps. Today, it powers various product features such as Messenger Contacts, Messenger Stories, Mobile Config etc.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Opportunity Store (Zippy)	Used for Lift and allows for a real-time join between opportunities (as well as some impression and click data) to conversions.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
orderdb (MySQL)	Orderdb contains historical transaction data, primarily from the time frame of 2009-2014, that is generally unused.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Payments (MySQL)	Payments teams are part of the Facebook Financial (F2) Org and power all Payments across all products and family of apps.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Perforce	supporting perforce users throughout the company.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Periplus (Viewstate)	Cache for the newsfeed data	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Pinpoint	Root cause analysis for metric movements (ex: Why did Facebook DAP go down between today and 7 days ago?)	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Poseidon	Poseidon is a storage service ontop of ZippyDB designed to store user side features.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Presence Storage Service	Power user presence (green dot) for Facebook, Instagram and Messenger users.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Puma	Puma is a a fully managed general purpose Stream Processing as a service solution (i.e. real-time data-processing).	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
PYMK Leaf Tailer	It is used to load realtime data into our leaf machines	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
RAIligator	Aggregation of call logs of different types by client session (i.e. one call connection) and call.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
		As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Readstate	Stores user interaction with the feed(i.e. all the stories user has seen)	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Ready at Dawn (M&A)	Ready at Dawn is a video game studio acquired by Facebook.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Recon DB	Recon DB is a data pipeline used for payment reconciliation	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.



Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Reliable Volumes	Reliable Volumes provides network attached data volumes to compute nodes, allowing services to scale compute and storage needs independently.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Research Home	To store long-term research artifacts submitted by UX Researchers for easy discovery.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
RocksDB (incl. Stargate)	RocksDB is a persistent key-value store library which can be used by anyone to store data on disk.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
RSC Storage	RSC Storage is a research / experimentation to do self-supervised learning on copy of production data	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Scribe	Scribe is a distributed, buffered, named pipe that serves as the entry point to most of the data ingestion pipelines at Facebook.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Scuba	The Scuba UI is a tool for interactively querying and visualizing data without writing SQL.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
SecureDumbo	Blob Storage System for storing backups of other systems in secure environment	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
ServiceFriend	a system to allow agents (humans and bots) to chat with customers	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
ServiceLab	Performance and Efficiency regression A/B testing platform.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
SFTP	SFTP is a service provided by the EPI team that allows Facebook to exchange data within the company and also with external vendors.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
ShadowDB	MyShadow is a service that captures all incoming traffic on a MySQL instance and replays it on another test MySQL instance in real time.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Sisyphus	Run batch data processing jobs over msgdb	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.



Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Sledge (Novi)	We store user's balance and transactions for the Novi wallet.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Sparse Inner Product / Feed Indexer	candidate generation for recommendation ranking system	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Sparta	Sparta is an application developer review tool meant to investigate cases where app devs could be misusing user data.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Stylus	Stylus is a general purpose Stream Processing library (i.e. real-time data-processing).	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Swift	Openstack Swift is an object storage service being provided by the EE Storage team. Currently it is aimed to be a low cost solution for archival workloads such as backups and data warehousing.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Synapse	Synapse is a distributed transactional storage system designed to give full composability across data structures, transactions, and storage encoding.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Tableau Server	Tableau is a business intelligence tool used to create charts, graphs, and other visualizations.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
TacoSRS	SRS (Single replica service) is a special variant of now deprecated TACO created as part of an effort to consolidate the non-replicated FQDB offering to core-data. SRS has two independent instances for feed tracking data and featuredb storage.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Tally	Tally can be seen as real-time counter service. It keeps statistics about streams of events, such as counts, unique counts, quantiles (percentiles), mean, and top elements, and it allows dynamic time window queries in real-time.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Tally on Sigma	Tally on Sigma is a real-time counter service for high volumes of data.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
TAO	TAO is a high-performance service for storing, caching, and querying the graph of FBObjects and associations.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	This is one of the various databases that powers the Social Graph. Facebook produced each Named Plaintiff's DYI file, which for each individual user represents the most complete and best compilation of data Facebook maintains associated with that user, and the best available compilation of the data about that user in the Social Graph, in a human readable form.
Threat Profile	Internal tool maintained by Global security to protect Facebook, executives and assets.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
TimelineDB	TimelineDB is a backend system that persists all actions by users and pages and indexes them chronologically.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
TRU	Transitive Resource Utilization (TRU) is a platform that leverages distributed tracing (Artillery) to collect service performance metrics and dependency context on a per-request basis across FB infrastructure.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
TW Storage	Tupperware (TW) is a managed deployment system for internal backend services.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
UDB	UDB is a backing database for TAO for user data.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	This is one of the various databases that powers the Social Graph. Facebook produced each Named Plaintiff's DYI file, which for each individual user represents the most complete and best compilation of data Facebook maintains associated with that user, and the best available compilation of the data about that user in the Social Graph, in a human readable form.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
unittestdb (MySQL)	UnitTestDB is a grouping of dbtypes that are used for testing.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Up2x	Up2x is a low-latency, high-availability, key-value storage system for user data.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
User Interest Profile (UIP)	Provide human readable understanding of user interests that are used to improve personalization in various user facing products.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
VertexDBProdNet	VertexDBProdNet is a tax application used for tax calculation for all FB products (Digital, Ads, Services and Hardware) globally.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.



Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Warm Storage	Warm Storage is a highly-distributed, highly scalable storage solution built to replace the large use cases of HDFS within Facebook.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
WhatsApp (Deletion Framework)	WhatsApp Messenger is a freeware, cross-platform and end-to-end encrypted instant messaging application for smartphones.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Wormhole	Wormhole is a publish-subscribe platform. It allows different groups at Facebook to (re-)use a well designed architecture to receive an ordered and reliable stream of data changes, such as user database (UDB) changes.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
XDB	XDB is Facebook's self-service relational database allocation system.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
XStream	XStream is Facebook's next generation streaming processing service.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
XWF-L	XWF-L is a platform which offers wifi as a product.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
ZippyDB	ZippyDB is a reliable, consistent, highly available, scalable key-value storage service.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	This is one of the various databases that powers the Social Graph. Facebook produced each Named Plaintiff's DYI file, which for each individual user represents the most complete and best compilation of data Facebook maintains associated with that user, and the best available compilation of the data about that user in the Social Graph, in a human readable form.
Muddler & HulkDB	Generating training data for instagram ML systems	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.

Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Double Vision	Record and serve IG user engagement history.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
IGML SIP	Embedding indexing system	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
AI Compliance Feature Store Features	Detect if any feature contains Ull and file the privacy-wave tasks for user verification/remediation	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
Mapillary	Public street-level imagery and map data platform	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.



Data System Name	High-Level Description of Data System	Explanation for Why Additional Information Is Not Readily Available	Explanation for Why Facebook Cannot Determine if the Named Plaintiffs' Data is Stored in the System
Fastdesk	Process WhatsApp user support requests	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.
RPM	RPM is a Code Management system owned by Core Systems.	As discussed in the Declaration of David Pope, Facebook has not previously compiled information responsive to the Special Master's Order—including descriptions of the "most common functions and purposes of the data system" and a comprehensive list of "the business units, divisions, or groups that use the data system"—in the course of its prior efforts to inventory its data systems. Facebook spent more than one year on this effort to inventory its systems, which required Mr. Pope's team to consult with hundreds of employees across dozens of departments in order to collect the information currently available, and collecting and vetting further detail about each system would require a similar effort that would be disproportionate to the needs of this case, as Plaintiffs' request to identify "all" user data is not relevant to any pending claims or defenses and the vast majority of such data could not readily be identified by user and/or produced in a format that would be unintelligible outside of Facebook.	The Special Master's Amended Order asks Facebook to determine whether each data system contains data related to a specific Named Plaintiff. The burdens associated with providing that information are not meaningfully different than identifying any specific data relating to a Named Plaintiff within each system, because—as discussed in the Ji and Pope Declarations—Facebook's data systems are generally not designed to allow for the identification and export of data related to a specific user and Facebook does not presently have an existing tool that would allow it to either identify or retrieve data related to a specific user from across all its data systems (the DYI tool retrieves data from a number of data systems, but not all). As a result, in order to determine what types of user data are contained within each data system and whether an individual user's data could be identified within that data, Facebook would need to undertake an analysis of each separate data system, how it is used, and whether it can be searched by user, which would require Facebook to repeat many of the steps undertaken by Mr. Pope's team over more than one year and present an ongoing moving target, given that the data in many of these data systems is not static.


# Exhibit K

GIBSON DUNN

Gibson, Dunn & Crutcher LLP

333 South Grand Avenue  
Los Angeles, CA 90071-3197  
Tel 213.229.7000  
www.gibsondunn.com

Deborah L. Stein  
Direct +1 213.229.7164  
Fax: +1 213.229.6164  
DStein@gibsondunn.com

January 27, 2022

VIA JAMS ACCESS

**HIGHLY CONFIDENTIAL - ATTORNEYS' EYES ONLY**

Re: *In re Facebook, Inc. Consumer Privacy User Profile Litigation*, 3:18-md-02843

Dear Special Master Garrie:

I write in response to your January 18, 2022 Supplemental Order Re: Facebook's Production of Plaintiff Data, in which you ordered Facebook to provide a "date certain" by which it could: (1) identify "[w]hich systems identified by Mr. Pope . . . contain discoverable data related" to the eight remaining Named Plaintiffs, Jason Ariciu, Bridgett Burk, Cheryl Senko, Jordan O'Hara, Samuel Armstrong, Steven Akins, Terry Fischer, and Tyler King, and (2) for those systems, "either produce th[at] . . . data from the system or articulate, with sufficient detail, the reason for it should not have to produce the Named Plaintiffs data from that system," and your supplemental January 25, 2022 Order.

Below we:

- i. Outline steps Facebook has taken to date to fulfill its discovery obligations with respect to data relating to Jason Ariciu, Bridgett Burk, Cheryl Senko, Jordan O'Hara, Samuel Armstrong, Steven Akins, Terry Fischer, and Tyler King;
- ii. Address the Special Master's request that Facebook identify which data systems contain discoverable data relating to Jason Ariciu, Bridgett Burk, Cheryl Senko, Jordan O'Hara, Samuel Armstrong, Steven Akins, Terry Fischer, and Tyler King and provide a narrowed list of the sources that are potentially-relevant to the Special Master's inquiry; and,
- iii. Describes Facebook's proposed next steps.

We respectfully remind the Special Master that discovery must be both relevant and proportional, considering, among other factors, "the importance of the discovery in resolving the issues, and whether the burden or expense of the proposed discovery outweighs its likely benefit." Fed. R. Civ. P. 26(b)(1). Courts have long recognized that the vast expansion of electronically stored information in the digital age has magnified the risk of discovery abuse: "[t]he information explosion of recent decades has greatly increased both the potential cost of wide-ranging discovery and the potential for discovery to be used as an instrument for delay or oppression." Fed. R. Civ. P. 26(b) advisory committee's note to 1993 amendment. Rule 26 was amended in 2015 to emphasize and mandate proportionality in discovery. "No longer is it good enough to hope that the information sought might lead to the discovery of admissible evidence. .

## GIBSON DUNN

Page 2 of 7

**HIGHLY CONFIDENTIAL – ATTORNEYS’ EYES ONLY**

. . . Instead, a party seeking discovery of relevant, non-privileged information must show, before anything else, that the discovery sought is proportional to the needs of the case.” *Gilead Scis., Inc. v. Merck & Co, Inc.*, 2016 WL 146574, at \*1 (N.D. Cal. Jan. 13, 2016).

As we discuss more below, the burdens associated with producing additional data relating to Jason Ariciu, Bridgett Burk, Cheryl Senko, Jordan O’Hara, Samuel Armstrong, Steven Akins, Terry Fischer, and Tyler King are not limited to extracting data and processing it for production. The mere process of investigating and searching numerous data systems (and multiple use cases within each of them) that are designed to provide a product to end-users—and not to be searched for individual users’ data—is itself extremely time consuming and burdensome. These burdens are far out of proportion to the needs of the case. This is particularly true given the volume of user data Facebook has produced to date and that conducting a scorched-earth search for duplicative data, or additional raw data that Facebook has confirmed is stored in a source not accessible to third parties, will not help Plaintiffs confirm what data was shared with third parties, which we understand to be the rationale for this exercise.

### **A. Facebook’s Efforts to Date.**

1. Facebook has complied with its discovery obligations to produce relevant user data related to the Named Plaintiffs and subsequent requests for information about the systems in which such data may be stored.

2. Facebook has produced nearly one million pages of data relating to current and former Named Plaintiffs in this action, including the individual user data relating to Jason Ariciu, Bridgett Burk, Cheryl Senko, Jordan O’Hara, Samuel Armstrong, Steven Akins, Terry Fischer, and Tyler King that is relevant to their non-stayed claims and does not present an undue burden on Facebook to search for, identify, collect, and produce. Facebook has provided the Special Master examples of the data it has produced for current and former Named Plaintiffs, which in some cases has exceeded 200,000 pages of data for a single Named Plaintiff. Facebook has also produced more than 80 GB of raw data memorializing Facebook’s transfer of data to third parties.

3. Facebook has developed a significant record demonstrating that the user data relevant to this case has been produced and that searching for additional user data relating to Jason Ariciu, Bridgett Burk, Cheryl Senko, Jordan O’Hara, Samuel Armstrong, Steven Akins, Terry Fischer, and Tyler King from its data systems would be unduly burdensome and disproportionate to the needs to this case. Specifically, Facebook has:

- a. Provided sworn testimony that Facebook makes user data available to third parties through application programming interfaces (“APIs”) (Anand Decl. ¶ 4);

## GIBSON DUNN

Page 3 of 7

**HIGHLY CONFIDENTIAL – ATTORNEYS’ EYES ONLY**

- b. Provided sworn testimony that APIs available to third parties “pull data exclusively from the Social Graph” (Anand Decl. ¶ 4);
- c. Identified the primary databases that support the Social Graph (*see* Ji Decl. ¶ 5);
- d. Provided sworn testimony that the DYI file “represents the most complete and best compilation of data Facebook maintains associated with [a] . . . user,” including—but not limited to—data from the Social Graph (Mitchel Decl. ¶ 5);
- e. Provided sworn testimony identifying the burdens of collecting user data beyond that reflected in the DYI file (*see* Ji Decl. ¶¶ 10-31);
- f. Provided sworn testimony that there is no existing tool—other than the DYI tool—that can be used to identify and retrieve data from a specific user across all of Facebook’s data systems and that endeavoring to engineer one would require a year of work by a team of engineers (Pope Decl. ¶ 10);
- g. Identified the relevant APIs through which third parties could access the categories of user data Plaintiffs have requested, identified the type of user data each API could be called for, and provided lists of the third parties that had access to them (*see* 4th Rog Responses);
- h. Provided sworn testimony that identifying additional data points about each of the 149 data systems disclosed to the Special Master, including the “types of user data” within each one, would likely take an additional year of work (Pope Decl. ¶ 9); and
- i. Provided a witness for over three hours of testimony to provide further information about the 149 data systems disclosed, which included testimony that even a single one of those data systems could have numerous use cases—and each may need to be investigated individually to answer the Special Master’s questions fully.

4. Plaintiffs have not made *any* showing that any of the above representations are incorrect or incomplete. While Plaintiffs have argued that they are “entitled to discovery to test Facebook’s assertions about what it made accessible to or shared with third parties” (Pls.’ Reply ISO MTC Named Plaintiff Data at 3), they have received that, in the form of millions of pages of documents, sworn testimony, the opportunity to take dozens of depositions, and 80 GB of data regarding Facebook’s transfer of data to third parties.

5. Duplicative data, or additional raw data that Facebook has confirmed is stored in a source not accessible to third parties, will not help Plaintiffs confirm what data was shared with third parties.

## GIBSON DUNN

Page 4 of 7

**HIGHLY CONFIDENTIAL – ATTORNEYS’ EYES ONLY**

6. At the Special Master’s request, Facebook has gone to great lengths to provide additional information about its data systems—which are some of the most technically advanced and complex systems in the world—through various submissions, declarations, and testimony. Also at the Special Master’s request, Facebook provided information about 149 different data systems (regardless of their relevance); it also provided a witness who appeared at a three-hour hearing with the Special Master to assist with the Special Master’s assessment of each of those data systems to the best of his ability.

7. Even with the benefit of all of this additional information, which goes far beyond the scope of proportional discovery to which they are entitled, Plaintiffs have not identified any specific information that is missing from Facebook’s productions or made any showing that Facebook’s representations demonstrating why its productions are complete are incorrect.

**B. The Special Master’s Request that Facebook Identify Which Specific Data Systems Contain Discoverable Data Relating to the Eight Remaining Named Plaintiffs.**

8. The Special Master asked Facebook to identify “[w]hich systems identified by Mr. Pope . . . contain discoverable data related” to the eight remaining Named Plaintiffs, Jason Ariciu, Bridgett Burk, Cheryl Senko, Jordan O’Hara, Samuel Armstrong, Steven Akins, Terry Fischer, and Tyler King.

9. As Facebook described in detail in the Pope declaration, Facebook cannot confirm with certainty that a particular data system contains individual information relating to a particular user—let alone data relevant to this litigation—without first conducting a complex analysis. As an initial matter, individual data systems may have multiple different use cases, all of which may store and use data differently and need to be investigated separately. 1/14/22 Hr’g Tr. at 87:13-23, 112:2-11. For each individual use case within each system, Facebook would need to assess whether the use case involves storage of individually-identifiable user data, determine whether the data sets used can be readily searched by user, and (if so) conduct a search of those data sets to determine whether data related to individual users exists in each system; assess the nature of the data and whether it is relevant and not duplicative of materials already produced; and determine whether that data can be extracted in a producible form. As the December 10 Ji Declaration made clear, this assessment and search could take months or years for a single data system. As Mr. Pope testified during the January 14 hearing, his team did not investigate what (if any) categories of user data are contained within each of the 149 systems he disclosed—or even whether each data system actually contains user data. 1/14/22 Hr’g Tr. at 8:12-9:8, 127:6-8 (“So the questionnaire will determine whether it may store user data, but not whether it does and what type it is.”). As relevant here, Mr. Pope’s team investigated only whether each system was technically capable of housing user data writ large and identified 149 such systems. *Id.*

10. Facebook has continued to investigate the 149 data systems identified in the Pope declaration and discussed at the January 14 hearing in order to narrow that list to systems that

## GIBSON DUNN

Page 5 of 7

**HIGHLY CONFIDENTIAL – ATTORNEYS’ EYES ONLY**

could *potentially* be relevant to the Special Master’s inquiry. Below, Facebook provides a narrowed (but still over-inclusive) list of 55 systems. This list excludes certain types of systems that Facebook understands are not relevant to the Special Master’s inquiry. Specifically, the list excludes:

- internal test systems,
- systems Facebook’s teams have declared not to store user data (even though as a technical matter, they are capable of storing such data),
- systems that do not serve the Facebook product and serve only other Meta products (such as WhatsApp),
- systems that do not store unique data, for example serving only as pipelines to transfer data from one system to another (such systems are designed only to move data and data cannot be read from them), and
- other systems from which Graph API cannot call data directly.

11. The remaining systems are as follows:

[REDACTED]

[REDACTED]

## GIBSON DUNN

Page 6 of 7

**HIGHLY CONFIDENTIAL – ATTORNEYS’ EYES ONLY**

[REDACTED]

[REDACTED]

12. Even within this narrowed list of 55 systems, certain systems may not [REDACTED]  
[REDACTED] For example, a system could store instances of [REDACTED]  
[REDACTED]

13. Data housed in Facebook’s systems can typically be [REDACTED]  
[REDACTED] through a complex series of connections within the Social Graph, which is powered by software code drawing on several different data systems and assets. In essence, the Social Graph brings together data from various systems and databases in a way that allows data to be connected so Facebook’s product can operate. Accordingly, [REDACTED]  
[REDACTED], and unless a system is connected into the Social Graph, and [REDACTED]  
[REDACTED]

14. The following subset of systems are connected to the Social Graph and therefore the data within them typically can be associated to [REDACTED]  
[REDACTED]

[REDACTED]

---

<sup>1</sup> Facebook notes that the DYI files contain IP addresses associated with the users’ log-ins.



GIBSON DUNN

Page 7 of 7

**HIGHLY CONFIDENTIAL – ATTORNEYS’ EYES ONLY**

**C. Further Response to the Special Master’s Inquiries.**

15. The Special Master asked for a date certain by which Facebook would identify the specific systems that contain data relating to Jason Ariciu, Bridgett Burk, Cheryl Senko, Jordan O’Hara, Samuel Armstrong, Steven Akins, Terry Fischer, and Tyler King and by which it would either produce the relevant data from those systems or explain why it should not have to do so.

16. As Facebook has explained numerous times, Facebook cannot determine with certainty if particular systems house “data” associated with specific users—leaving aside whether such data could possibly be relevant here and has not already been produced—without conducting an extremely time-consuming and burdensome investigation into individual data systems and each of their use cases. Facebook’s systems are engineered to provide its website and apps to users, not for production in litigation of all data associated with a specific user. Existing tools are not built to retrieve all data associated with a single user, and creating such a tool requires significant engineering resources on top of coordination with the teams that operate each system, as well as the product teams that build the tools and products that use each system.

17. Facebook has now narrowed the 149 systems it disclosed originally to 12 systems that could be *potentially* relevant to the Special Master’s inquiry, so that it is able to provide the Special Master more concrete information about each.

18. Facebook suggests that it work with the Special Master to identify three systems from the list of 12 above, and within 14 days, Facebook will provide a detailed analysis with evidentiary support regarding the steps that would be necessary for those particular systems to identify whether they house data Facebook has not already produced relating to Jason Ariciu, Bridgett Burk, Cheryl Senko, Jordan O’Hara, Samuel Armstrong, Steven Akins, Terry Fischer, and Tyler King, and whether it is possible to extract that data in a producible format. Facebook expects that even this narrowed exercise would be an extremely burdensome and disproportionate effort.

As always, we are ready and available to discuss this issue further with the Special Master.

Respectfully submitted,



cc: Lesley E. Weaver  
Derek W. Loeser

# Exhibit L

HIGHLY CONFIDENTIAL

Page 1

JAMS

---o0o---

IN RE: FACEBOOK INC., )  
CONSUMER PRIVACY USER )  
PROFILE LITIGATION )  
\_\_\_\_\_ )

HIGHLY CONFIDENTIAL

HELD VIA ZOOM

REPORTER'S TRANSCRIPT OF PROCEEDINGS:

JAMS Special Master Hearing

Thursday, February 17, 2022

REPORTED BY:

Katy E. Schmidt

RPR, RMR, CRR, CSR 13096

Job No.: 5096673

HIGHLY CONFIDENTIAL

Page 2

JAMS

---o0o---

IN RE: FACEBOOK INC., )  
CONSUMER PRIVACY USER )  
PROFILE LITIGATION )  
\_\_\_\_\_ )

JAMS Special Master Hearing on Thursday,  
February 17, 2022, at 1:07 p.m., virtually before  
Kathryn E. Schmidt, RPR, RMR, CRR, CSR 13096.

HIGHLY CONFIDENTIAL

Page 3

1 APPEARANCES:

2 For The Plaintiffs:

3 (Appeared via Zoom)

4 KELLER ROHRBACK LLP

5 BY: DEREK LOESER, Esq.

6 BY: CARI LAUFENBERG, Esq.

7 BY: BENJAMIN GOULD, Esq.

8 BY: CHRIS SPRINGER, Esq.

9 1201 Third Avenue, Suite 3200

10 Seattle, Washington 98101-3052

11 dloeser@kellerrohrback.com

12 BLEICHMAR FONTI & AULD LLP

13 BY: LESLEY WEAVER, Esq.

14 BY: MATT MELAMED, Esq.

15 BY: ANNE DAVIS, Esq.

16 555 12th Street, Suite 1600

17 Oakland, California 994607

18 415.445.4003

19 lweaver@bfalaw.com

20 For The Defendants:

21 (Appeared via Zoom)

22 GIBSON DUNN & CRUTCHER LLP

23 BY: RUSS FALCONER, Esq.

24 BY: MARTIE KUTSCHER CLARK, Esq.

25 BY: LAURA MUMM, Esq.

BY: IAN CHEN, Esq.

BY: ROSE RING, Esq.

BY: COLIN DAVIS, Esq.

2001 Ross Avenue, Suite 2100

Dallas, Texas 75201

214.698.3170

rfalconer@gibsondunn.com

---o0o---

HIGHLY CONFIDENTIAL

Page 4

1 LOS ANGELES, CALIFORNIA

2 THURSDAY, FEBRUARY 17, 2022

3 ---o0o---

4 SPECIAL MASTER GARRIE: All right. All right.  
5 So we'll go on the record.

6 Okay. This is Special Master Garrie. We have  
7 a hearing today to discuss named plaintiff's data.

8 I asked Facebook to provide a series of  
9 witnesses to get further clarification as the objective  
10 of today is to determine which repositories are  
11 appropriate from the initial 149 repositories that were  
12 identified by Facebook as potentially having Facebook  
13 data, that was subsequently narrowed down to 55  
14 potential systems as potentially having -- may -- I  
15 underline, may have user data in it.

16 The goal of today is to determine what -- for  
17 me to get technical information to figure out what is  
18 reasonable for Facebook to look for the named plaintiff  
19 data.

20 For purposes of convenience today, named  
21 plaintiff data and user data are one in the same.  
22 They're not -- and it's not aggregated. It is user  
23 specific data. It is not anonymized data. It is not  
24 de-anonymized data. It is their data. It is user data.

25 Okay?

HIGHLY CONFIDENTIAL

Page 5

1           So with that sort of definition, I will use  
2           them interchangeably, user data and named plaintiffs,  
3           but to me they're users.

4           The goal of today is to determine what is  
5           reasonable and what -- where -- what systems have what  
6           data that relates to the named plaintiffs pursuant to  
7           the Court's definition.

8           With that in mind, I wanted to read out for  
9           really the Facebook engineers that are here. I'm glad  
10          you guys are here, because I feel like I can finally  
11          just get -- we can make -- we can make progress.

12          I just want to read the definition that the  
13          judge has defined as what is discoverable user data  
14          issue.

15          Okay?

16          Because it's not -- it's a defined definition  
17          and I think it's worth reading for everybody.

18          So the first is "Data collected from a user's  
19          on-platform activity."

20          Okay?

21          So it's "Data collected from a user's  
22          on-platform activity." That's No. 1.

23          No. 2, "Data obtained from third parties  
24          regarding a user's off-platform activities."

25          Now, there's no definition as to what

HIGHLY CONFIDENTIAL

Page 6

1 off-platform activities are or anything of the like.

2 So with that in mind, we then have No. 3,  
3 "Data inferred from a user's on-or-off platform  
4 activity."

5 So if it's obtained from third parties.

6 And then group No. 3 is "Data inferred," which  
7 again is not defined -- but "inferred from a user's  
8 on-or-off platform activities."

9 I take all platform activity to be activity  
10 that is not being done within the je ne sais quoi of  
11 Facebook, whether it's -- I don't -- there's lots of  
12 assets, so off your platform.

13 Okay?

14 With that in mind, that's the general set of  
15 definitions.

16 Okay. So with that in mind, I have a set of  
17 questions, and I assume I can start with Mr. Zarashaw.

18 There are two important things you have to  
19 remember, and I know both of you have been -- had the  
20 privilege of meeting the lawyers here from both sides.

21 The attorney-client privilege is the only  
22 thing that really counts today; right? If there's  
23 something that you talked about with your lawyers, say  
24 I -- your lawyers can object, saying it's -- or you can  
25 say "I discussed this with my lawyers," whatever. I



HIGHLY CONFIDENTIAL

Page 7

1 think for Facebook it will be Counsel Falconer making --  
2 or note -- assisting you as necessary, and it will be --  
3 well, plaintiffs won't object to attorney-client  
4 privilege, so --

5 But the goal today isn't to learn anything  
6 about -- anything but systems and data.

7 Okay?

8 I'm not -- I mean, I appreciate the underlying  
9 case and the zealous advocacy, but the truth is what I'm  
10 here to do today is to learn and determine about your  
11 systems. Like because you've disclosed there are 55 of  
12 them that may have plaintiff data, and I think that may  
13 be a little excessive to require Facebook to look  
14 through all 55 systems for user data, given these  
15 definitions.

16 So then with that in mind, I'm going to kick  
17 it off --

18 Or are there any questions from Facebook?

19 MR. FALCONER: One or maybe two.

20 Is it okay with the Special Master if Mr. Elia  
21 and Mr. Zarashaw -- there is a list of questions that we  
22 answered on February 9th that we -- public information  
23 filed into the case.

24 Is it acceptable to the Special Master if the  
25 witnesses want to refer back to those written answers

HIGHLY CONFIDENTIAL

Page 8

1 during the hearing today?

2 SPECIAL MASTER GARRIE: Yeah.

3 MR. FALCONER: Okay.

4 SPECIAL MASTER GARRIE: Of course.

5 MR. FALCONER: And --

6 SPECIAL MASTER GARRIE: Go ahead.

7 MR. FALCONER: And then just for the  
8 Special Master's benefit, if this is helpful,  
9 Mr. Zarashaw is here specifically with expertise on if  
10 there was a API misused topic that the Special Master  
11 indicated he was interested in. And that's -- that's  
12 Mr. Zarashaw's kind of area of knowledge and expertise.

13 The other topics that the Special Master had  
14 indicated interest in, Mr. Elia is the person most  
15 knowledgeable on those topics, if that helps -- just if  
16 that helps with kind of organization and flow.

17 SPECIAL MASTER GARRIE: It certainly does.  
18 I'm just working on my mute and un-mute skills. It  
19 certainly does.

20 Yeah. And if either one of the technical  
21 witnesses don't feel comfortable answering a question,  
22 I get that there are millions of lines of source code.  
23 I get there are thousands of engineers. I fully  
24 understand that being asked a question on the spot may  
25 not result in an exact answer.

## HIGHLY CONFIDENTIAL

Page 9

1           The idea is just to move it forward so that I  
2           can at least put together tangible understanding of what  
3           is appropriate and not appropriate or reasonable and not  
4           reasonable in the context of the named plaintiff data.

5           So -- okay. So API, so I guess Mr. -- oh,  
6           sorry.

7           Did plaintiffs have any questions?

8           MS. WEAVER: No. Not at this time. Thank  
9           you.

10          SPECIAL MASTER GARRIE: Okay. So,  
11          Mr. Zarashaw, I have some questions about APIs.

12                 Before we get into that, I guess my first  
13                 question is what types of API's interaction took place,  
14                 specifically applications interacting with user data  
15                 versus the bulk API integrations with partners; right?  
16                 Because there's different ways and there's these bulk  
17                 APIs and then the other way.

18                 What types of sort of API interactions sort of  
19                 played out, or possible?

20                 MR. EUGENE ZARASHAW: So I'm not familiar with  
21                 a distinction between individual versus bulk API.  
22                 However, there is a distinction between server to server  
23                 API calls made by a third-party app from their secured  
24                 servers to our servers versus calls made by an app in  
25                 the mobile device or in the web browser controlled by a

HIGHLY CONFIDENTIAL

Page 10

1 user to our servers on behalf of an app.

2 Is that the distinction?

3 SPECIAL MASTER GARRIE: Yeah. So that --  
4 yeah. So then that's how they're organized?

5 So -- because there's been a lot of terms  
6 batted around about the bulk API integration from  
7 partners, and I didn't fully understand what that meant.  
8 So I'm glad that I'm not fully --

9 MR. EUGENE ZARASHAW: If I may add to that,  
10 the APIs are not quite organized that way. There's  
11 overlap in APIs that can be called only from a mobile  
12 client, slash, specified only server to server or both,  
13 depending -- and the distinction tends to be in the  
14 access token used to access the API rather than API  
15 itself.

16 SPECIAL MASTER GARRIE: Okay. Got it.

17 MR. EUGENE ZARASHAW: Would an example help?

18 SPECIAL MASTER GARRIE: Yeah.

19 MR. EUGENE ZARASHAW: So, for example, any  
20 time a user is running the Facebook app or a third-party  
21 app on a mobile device, the access token on that mobile  
22 device inside the app only should have access to data  
23 accessible to the user normally, intersected with the  
24 data that that app should have access to.

25 So a subset of the two.

## HIGHLY CONFIDENTIAL

Page 11

1           That -- which means if there's any data not  
2       accessible to the user that owns that access token, then  
3       that access token cannot retrieve that data.

4           So as an example, if you and I are not friends  
5       on Facebook and have -- I cannot see your data on  
6       Facebook, my access token cannot be used to fetch data  
7       about you.

8           By comparison, there's the concept of an  
9       application access token where in a server-to-server  
10      call, if you and I have both installed the application  
11      and given the application permission to look at our  
12      data, that application access token can be used to fetch  
13      data about both of us on behalf of just the application.

14           SPECIAL MASTER GARRIE: So then what if -- is  
15      that true for mobile as well then?

16           MR. EUGENE ZARASHAW: On mobile we only use  
17      the first type of access tokens, user access tokens,  
18      because it would be unsecure to allow the server side  
19      access token out of the control of the app developer's  
20      servers.

21           SPECIAL MASTER GARRIE: Okay. So then does  
22      Facebook's APIs, based on social graph and allow third  
23      parties to access information on the user's friends  
24      through that same system -- well, that is my first  
25      question, and then using the system you just described.

HIGHLY CONFIDENTIAL

Page 12

1 MR. EUGENE ZARASHAW: So using the system I  
2 just described, there would be two ways to access  
3 information about a user's friends. If we're talking  
4 about the mobile scenario --

5 SPECIAL MASTER GARRIE: Yeah, mobile.

6 MR. EUGENE ZARASHAW: In the mobile -- then in  
7 the mobile scenario, the access token could be used to  
8 call an API to get a user's friends who, after the  
9 platform 2.0 changes, only return the friends who have  
10 also installed the app and granted the app permissions.

11 SPECIAL MASTER GARRIE: Yeah.

12 MR. EUGENE ZARASHAW: So it would return a  
13 sub -- the intersection of the friends the apps has  
14 access to, and the friends that the user has access to.

15 SPECIAL MASTER GARRIE: And then in  
16 Version 1.0?

17 MR. EUGENE ZARASHAW: It would have returned  
18 all of a user's friends if that friend had granted  
19 friend permissions.

20 SPECIAL MASTER GARRIE: And that was in the  
21 interrogatory's full response so that makes sense.

22 So then -- so then the Facebook's API is based  
23 on the social graph, so then the answer is yes. Okay.

24 So then -- because one of the things -- so  
25 there's another been a lot of talk of social graph;

HIGHLY CONFIDENTIAL

Page 13

1 right? And social graph is a -- I'll just leave it  
2 there. And it allows third parties basically to  
3 continue access --

4 Does the -- does Facebook's APIs based on the  
5 social graph really allow third parties then to continue  
6 accessing friends of friends of friends, like sort of a  
7 Russian nesting doll, for lack of a better term?

8 MR. EUGENE ZARASHAW: They did not unless the  
9 app had access to those specific friends.

10 So there's always the intersection of what  
11 should the app be able to see and what should the user  
12 be able to see? The app could only see information  
13 about people who had installed the app, or in 1.0, only  
14 information about people who are friends with someone  
15 who installed the app when acting on behalf of the  
16 person that installed the app.

17 So you could not just -- you could not find  
18 data about unrelated people to whom the app should not  
19 have access to.

20 SPECIAL MASTER GARRIE: Okay. All right.  
21 That makes sense.

22 So then one of the things I want to  
23 understand, because there's all this talk of all these  
24 different systems, but one thing I'm trying to  
25 understand is could then a third-party application

HIGHLY CONFIDENTIAL

Page 14

1 present a user with a user interface that mimics a  
2 Facebook log-in page and collect their credentials, and  
3 then log into the Facebook application on its own?

4 So I can walk you through what I'm talking  
5 about, but I think you get the gist of it. I mean...

6 MR. EUGENE ZARASHAW: I think it might help  
7 walking through, actually, if you don't mind.

8 SPECIAL MASTER GARRIE: Yeah.

9 So a third-party application, an app developer  
10 builds something, and presents a user with a user  
11 interface that mimicks Facebook's log-in page.

12 Okay?

13 So they build an app. They think they're  
14 logging in. They collect the credentials.

15 Okay?

16 So the app collects the credentials.

17 Then the app logs into the Facebook  
18 application on its own.

19 Is that -- is that possible?

20 MR. EUGENE ZARASHAW: So it was technically  
21 possible. It was something that we only allowed to  
22 specific apps under a contract. It was called the  
23 log-in permission.

24 SPECIAL MASTER GARRIE: That's it.

25 MR. EUGENE ZARASHAW: And it would have -- the



HIGHLY CONFIDENTIAL

Page 15

1 app would have to be permitted to exchange a user name  
2 and password for a Facebook access to open to log in  
3 this way.

4 And only apps who -- that had their log-in  
5 permission were able to do it.

6 This was used in cases where usually due to  
7 device constraints, it was impossible to launch a full  
8 web browser and log into Facebook, such as an embedded  
9 device.

10 SPECIAL MASTER GARRIE: Or a Nokia flip phone  
11 or something?

12 MR. EUGENE ZARASHAW: (Nodding.)

13 SPECIAL MASTER GARRIE: So then -- you said  
14 there was an agreement. And then they could log in, but  
15 they would have a different level of access -- the app  
16 would have a different level of access because they  
17 would be -- have the user's credentials they're logging  
18 in, not the other way; right?

19 MR. EUGENE ZARASHAW: Yes.

20 SPECIAL MASTER GARRIE: Okay. And then today  
21 you've sort of stopped that, though. As I understand,  
22 it has evolved, your user interface, to prevent -- to  
23 limit that ability today.

24 MR. EUGENE ZARASHAW: May I ask Steven to  
25 answer this question, just because I'm a bit out of date

HIGHLY CONFIDENTIAL

Page 16

1 on the current state?

2 SPECIAL MASTER GARRIE: Steven, hello again.

3 I hope you had a good evening at least.

4 MR. STEVEN ELIA: Yeah.

5 I'm not sure whether there's currently any of  
6 those types of special log-in whitelists that remain.

7 SPECIAL MASTER GARRIE: But there would be a  
8 contract or agreement or something there; right?

9 MR. STEVEN ELIA: That's right.

10 SPECIAL MASTER GARRIE: Okay. That's fine.

11 Then --

12 MS. WEAVER: Special Master Garrie -- I  
13 apologize. But I think it would be helpful for us to  
14 know time frame, and I think you're going there, and I  
15 didn't know when to ask.

16 But when did the system operate that they just  
17 described? And maybe you're about to get there and I  
18 apologize.

19 SPECIAL MASTER GARRIE: I was literally --

20 MS. WEAVER: I've been getting texts so I  
21 apologize. Okay.

22 SPECIAL MASTER GARRIE: Here's what I'll do.

23 Plaintiffs could write down all of their  
24 questions. When I finish a section, plaintiffs can ask  
25 their questions in case I forget anything, but I'm

HIGHLY CONFIDENTIAL

Page 17

1 really -- this is more for me than anybody else.

2 So my question is with regards to the timeline  
3 for this -- what I just described; right? Right?

4 So at some point I know -- well, if you don't  
5 have it, it's completely acceptable; right?

6 Do you know when this feature set -- what I  
7 just described was restricted or removed from the  
8 Facebook world, I guess? Platform? Sorry.

9 MR. EUGENE ZARASHAW: I do not.

10 SPECIAL MASTER GARRIE: Steven?

11 MR. STEVEN ELIA: I believe there were still  
12 uses of it in 2018. Since then, I'm not sure.

13 SPECIAL MASTER GARRIE: Okay. That's fine.  
14 Better -- "I don't know" is completely an acceptable  
15 answer and it's something we can follow up on.

16 My next question is then: Based on the  
17 scenario I just described, can a third party log the  
18 user into GAE?

19 MR. EUGENE ZARASHAW: I'm sorry. But what is  
20 a GAE?

21 SPECIAL MASTER GARRIE: Sorry. It's the graph  
22 application -- you guys have all -- sorry. It's the  
23 graphical --

24 MR. STEVEN ELIA: Graph API Explorer?

25 SPECIAL MASTER GARRIE: Yeah. Thank you. The

HIGHLY CONFIDENTIAL

Page 18

1 Graph API Explorer. I wrote all these down. I'm just  
2 trying to keep track. There's a lot of APIs you guys  
3 have. The Graph API Explorer.

4 Could a third party log the user into the  
5 Graph API Explorer? Because that's -- in the scenario  
6 I just described?

7 MR. EUGENE ZARASHAW: Yes. With a user name  
8 and password and using Graph API Explorer, they could.

9 SPECIAL MASTER GARRIE: Yeah. They --

10 MR. EUGENE ZARASHAW: But their access would  
11 still be limited to only things the user would have  
12 access to.

13 SPECIAL MASTER GARRIE: Right.

14 But that's different according -- at least my  
15 reading of your API and the libraries, whatever, that's  
16 still a different level permission than the third-party  
17 app developer would be naturally granted.

18 MR. EUGENE ZARASHAW: It grants access to a  
19 number of things that would normally require app review  
20 and it would violate our policies if they were to do  
21 this, by the way.

22 SPECIAL MASTER GARRIE: No, no --

23 MR. EUGENE ZARASHAW: It no way bypasses the  
24 controls on what data should a user be able to see.

25 So there -- there's no app ID under which you

HIGHLY CONFIDENTIAL

Page 19

1 would be able to get more information with the access  
2 token than the user would normally have by just using  
3 the Facebook site, which if you had the users use their  
4 password, you could also just go to the Facebook site  
5 without the API at all and get the same data.

6 SPECIAL MASTER GARRIE: Right.

7 But -- exactly. But if you're a third-party  
8 app developer, you'd have to go through the -- and I get  
9 it violates your policies. I'm not saying anything of  
10 the sort.

11 But this is way to get a -- that would allow  
12 an app developer to circumvent the need to get the whole  
13 app -- I mean, I've heard a lot of testimony about your  
14 app developer policy -- that whole -- I get it is my --  
15 it allows them to circumvent it.

16 But I just -- so then -- so then I guess  
17 from my perspective, did Face- -- so then -- okay.  
18 All right.

19 So then the next question I have -- sorry.  
20 I'm just moving things around because I thought there  
21 would be more witnesses, so let me just...

22 API...

23 I have more API questions but let me -- I have  
24 a couple questions about data collection.

25 So are you the right person for third-party

HIGHLY CONFIDENTIAL

Page 20

1 integration APIs, Mr. Zarashaw?

2 MR. EUGENE ZARASHAW: I think I would have a  
3 subset of information.

4 SPECIAL MASTER GARRIE: Okay. Then I'll just  
5 continue down the path that I have. I'll just --

6 Let me rethink this whole plaintiff question  
7 strategy.

8 I think plaintiffs, if you have specific  
9 questions, I will hear them and then determine whether  
10 the witness needs to answer them.

11 So do you have any questions at this point?

12 MS. WEAVER: Just more specificity as to when  
13 in 2018.

14 SPECIAL MASTER GARRIE: Oh, that was simple.  
15 He did his best to give his answer and we can follow up.  
16 So I'm not going to ask.

17 All right. All right. So then activity on  
18 data collection, that is something just for Facebook to  
19 understand that I do think that should -- clarity should  
20 be provided as to when that access -- well, it doesn't  
21 really relate -- well, I've gotta think about it. Let  
22 me just keep going here.

23 Activity and data collection. So -- I  
24 believe, and I just want to verify, Facebook collects  
25 user activity, like page views, clicks, ad impressions,

HIGHLY CONFIDENTIAL

Page 21

1 et cetera; right?

2 Is that right, Mr. Zarashaw?

3 MR. EUGENE ZARASHAW: We collect data of user  
4 interaction with our site and app.

5 SPECIAL MASTER GARRIE: And do you collect  
6 data like page views?

7 MR. EUGENE ZARASHAW: Yes.

8 SPECIAL MASTER GARRIE: Mouse clicks?

9 MR. EUGENE ZARASHAW: Yes. Depending on  
10 which -- what the clicks are on.

11 SPECIAL MASTER GARRIE: Yeah. So if I click a  
12 like, do you capture that?

13 MR. EUGENE ZARASHAW: Exactly. If you click a  
14 link, yes. If you click somewhere randomly on the page,  
15 I don't believe we capture it.

16 SPECIAL MASTER GARRIE: So then I want to ask  
17 about mobile but I feel like that's a whole different  
18 animal and I don't know if I want to get to it now or  
19 wait.

20 But do you collect data for mobile, like --  
21 you know, there's all this data that mobile devices  
22 have, like your Bluetooth signal, your wifi network. I  
23 mean, theoretically, when they log in, as a developer, I  
24 recognize that information is accessible.

25 Does Facebook collect that information, like,

## HIGHLY CONFIDENTIAL

Page 22

1     you know, nearby wifi access points -- about the user;  
2     right? -- GPS location, whatever they connect with;  
3     right? I mean, do you guys collect that from the mobile  
4     side of the house per user?

5             MR. EUGENE ZARASHAW: I don't know what we  
6     collect on the mobile side.

7             SPECIAL MASTER GARRIE: Okay. So like time  
8     zone, language, cookie data.

9             MR. EUGENE ZARASHAW: I apologize. I've not  
10    spent time on the mobile side of the house so I don't  
11    know what we collect there.

12            SPECIAL MASTER GARRIE: All right. Well,  
13    let's talk about on the nonmobile side of the house.

14            What about like cookie data, like, that are  
15    stored on the computer that isn't yours, like cookie IDs  
16    and settings?

17            MR. EUGENE ZARASHAW: We only get cookies for  
18    our own domain sent to us as part of the HTTP requests.  
19    So we are not able to collect cookies outside of the  
20    Facebook domain or other domains.

21            SPECIAL MASTER GARRIE: Yeah. Like Facebook  
22    domains, like Instagram -- like, it's everything.

23            When you say Facebook domain, you're referring  
24    to the Facebook universe; right?

25            MR. EUGENE ZARASHAW: I can confidently say we



## HIGHLY CONFIDENTIAL

Page 23

1 do not collect data from other domains that we don't  
2 have control over. I do not know what exactly we do  
3 collect from the domains we do have control over.

4 SPECIAL MASTER GARRIE: Okay. And then like  
5 obviously -- so then one thing that I would want to know  
6 from Facebook is if you collect mobile phone numbers,  
7 IP addresses, connection speed, and other things tied to  
8 the mobile, accessing Facebook through the mobile side  
9 by user; right? Not anonymized. Not anything else.  
10 For the user; right?

11 And if -- Mr. Zarashaw -- or just Eugene and  
12 Steven -- Steven, did you by chance know the answer to  
13 my question?

14 MR. STEVEN ELIA: I don't know specific  
15 instances. I know at certain points we've had offerings  
16 in our app to, for example, import your contacts on your  
17 phone to be able to find friends that you're not yet  
18 connected to on Facebook.

19 SPECIAL MASTER GARRIE: From the mobile  
20 device; right? So when you connect the mobile  
21 connectivity, is that associated with the user, and  
22 you're obviously granting permission for other things?

23 MR. STEVEN ELIA: I don't know a ton about  
24 specifically about how that works. I haven't worked on  
25 the mobile side. But I imagine that would read from

HIGHLY CONFIDENTIAL

Page 24

1 your, for example, phone book on your phone.

2 MR. FALCONER: Mr. Garrie, if it's helpful,  
3 we've got some documents that address this question that  
4 we could submit for the Special Master, too, if that  
5 would be helpful.

6 SPECIAL MASTER GARRIE: Oh, yeah. That's  
7 fine. I don't need it -- if you have the data to answer  
8 the question, you can submit it. It's fine. I just --  
9 there were so many documents and so much data that -- I  
10 mean, I did read your interrogatories 1, 2, 3, 4 -- I  
11 mean, there were a lot of them, and a lot of data, so I  
12 may have missed it. So if you could flag it, that would  
13 be helpful for me.

14 So then with that in mind, right, where is the  
15 user data that's collected stored? The user activity  
16 data that we just talked -- forget mobile. Just focus  
17 on what we know. Where is this user activity collected  
18 stored?

19 MR. EUGENE ZARASHAW: It's in a very wide  
20 variety of systems. I think the most helpful way to  
21 partition it might be to look at data that will at some  
22 point in the future expect it to be read again versus  
23 data that is being written and likely will not be read  
24 any time soon.

25 So I believe we've referred to it as the

HIGHLY CONFIDENTIAL

Page 25

1 social graph for the entire set of data we store that  
2 might also be used by the product again.

3 And then there are additional types of offline  
4 storage where we do things, like logging for every kind  
5 of reason, performance, monitoring, stability, just  
6 backup, and so on. And a lot of that may end up in  
7 other systems that can then be read from production.

8 SPECIAL MASTER GARRIE: So --

9 MR. EUGENE ZARASHAW: It could be read by the  
10 product.

11 SPECIAL MASTER GARRIE: Read by the product.

12 So this is the question; right? There are  
13 these 55 systems that have been identified. We can go  
14 through each one to figure this out or -- and they've  
15 identified a subset of 12. I get that's a social graph.  
16 Okay. Great. We can then say that we're down now to  
17 43 systems.

18 Okay?

19 There's still 43 systems that may have  
20 plaintiff data when the user activity is stored; right?  
21 Where is a user -- I mean, it would be good to have with  
22 some particularity where the actual user data that is  
23 collected by Facebook -- all of it; right?

24 Because there's a lot of data you collect when  
25 someone logs in. They don't -- the computer gives you;

HIGHLY CONFIDENTIAL

Page 26

1 right? They don't have an idea what their IP address  
2 is. They have no idea what version of their browser  
3 they're using. They have no idea what version of their  
4 operating system they're using.

5 But I assume you, like every other web-based  
6 application or platform, check these things, you know,  
7 part for security, part for functionality, part for the  
8 user experience; right?

9 Someone has a Mac that's four years old and  
10 you try to show them some things on the web, it will  
11 break the computer, and other --

12 You know, so you check these things and you  
13 store this and associate it with the user.

14 So what I'm interested in is what data you --  
15 but I get you don't keep all of that.

16 Okay?

17 But let's be clear there's user data, like  
18 page use.

19 Okay?

20 Let's focus on page use.

21 Where does the data for page use go? Like  
22 'cause I assume you track how long they look at it, they  
23 scroll, the whole nine yards.

24 MR. EUGENE ZARASHAW: The last I touched this  
25 in around 2015 or 2016, the primary repository for --

HIGHLY CONFIDENTIAL

Page 27

1 provide once and don't read until in bulk or during  
2 analysis data would have been in Hive. And there's  
3 quite possibly other specific systems it ended up in for  
4 data analysis and so on.

5 But Hive was the general repository for  
6 analytics data when there was no need for the product to  
7 ever read it again.

8 SPECIAL MASTER GARRIE: Got it.

9 And then I'm looking for this letter to  
10 show -- I'll share my screen in a second to go through  
11 the 50 -- because -- but here's the problem: Of the 43  
12 remaining systems, only one of them is Hive.

13 MR. EUGENE ZARASHAW: I commiserate with  
14 the --

15 SPECIAL MASTER GARRIE: There are 42 other  
16 systems that I'm trying to understand may possibly have  
17 user data.

18 MR. EUGENE ZARASHAW: I commiserate with the  
19 problem because looking at the list, I don't recognize a  
20 number of those names either.

21 SPECIAL MASTER GARRIE: Well, good. I was  
22 worried because I was like searching the data to make --  
23 that's already been provided to see if it existed.

24 So then --

25 MR. EUGENE ZARASHAW: I think the clearest

HIGHLY CONFIDENTIAL

Page 28

1 distinction, if we're looking for -- understanding what  
2 data then left Facebook again is --

3 SPECIAL MASTER GARRIE: We'll get to that. I  
4 have a whole other set of questions of what left, came  
5 back, received, and all of that. I have a whole set of  
6 questions about that.

7 So then where does impressions go? Like how  
8 many times I reload or load a page? Like, you know,  
9 you put -- I'm sure you track it for ads but also just  
10 impressions; right? Page impressions or ad -- ad  
11 impressions and page impressions are different; right?  
12 They're different units, so to speak?

13 MR. EUGENE ZARASHAW: So you've worked on  
14 systems -- you've worked on similar systems before where  
15 there are task specific storage systems.

16 SPECIAL MASTER GARRIE: Yeah.

17 MR. EUGENE ZARASHAW: So we have a -- we have  
18 a general set of [REDACTED] that happens in ads into Hive  
19 and paid and nonpaid impressions as well goes into Hive.

20 Because ads is a complex system in of its own,  
21 it also has ad specific systems that also get the same  
22 [REDACTED] to them and potentially use different  
23 sets of the data for more [REDACTED] approaches.

24 SPECIAL MASTER GARRIE: So then if a user  
25 clicks on an ad, right, and they engage with an actual

HIGHLY CONFIDENTIAL

Page 29

1 ad unit, you capture that and that ends up in the DIY  
2 (sic).

3 Like what I'm trying to understand is users  
4 and ad -- like page impressions or whatever, people are  
5 getting paid somehow.

6 So like how does the user relationship to,  
7 just as an example, the ad -- they view a page and they  
8 see an ad. How does that data get stored and captured  
9 and associated with the user so you don't show the same  
10 user the same ad 20 million times?

11 MR. EUGENE ZARASHAW: I'm not sure that we  
12 don't show the same user the same ad more than once.

13 I'm not sure exactly how we store it -- or I  
14 should say I'm not sure of all of the ways we store it.

15 But, for example, in case of an ad being shown  
16 to a person, [REDACTED]

[REDACTED]

[REDACTED]

19 [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

24 And I am not a deep expert so I'm probably  
25 missing other systems that it also could get written to.

HIGHLY CONFIDENTIAL

Page 30

1           SPECIAL MASTER GARRIE: What I'm trying to  
2 figure out is there's two systems, 1 and 2; right?  
3 Ad market DB? And I'm trying to figure out if you have  
4 user data relationship stored for the ad unit about what  
5 the user saw or is it duplicative? That's the first  
6 question. And then -- for systems 1 and 2.

7           And then my next very obvious question, is it  
8 duplicative of what you already put in elsewhere?

9           Because one of the things that nobody has told  
10 me yet is what is duplicative of the DIY file versus  
11 what is unique that's not in the DIY file that you can  
12 get from the systems?

13           So what I'm trying to figure out is, you know,  
14 if data is collected about user activity and it's stored  
15 elsewhere besides the Hive, where -- and those  
16 12 systems, where else is the data stored? Like Tableau  
17 servers, like, I store stuff on planet earth. I mean,  
18 it's big.

19           Is there any more -- like, do you have any  
20 idea where that information is actually stored today or  
21 was stored? Because these are huge systems; right?

22           Tableau is like they paid billions of dollars  
23 for it. It is a huge system in itself; right? I mean,  
24 there has to be some rhythm to where you're storing this  
25 information.



HIGHLY CONFIDENTIAL

Page 31

1 MR. EUGENE ZARASHAW: Not all of these systems  
2 are that huge.

3 So I think part of the rhythm would be that  
4 some of the systems are designed to work with much  
5 narrower sets of information, such as not having user  
6 cardinality data sizes. And instead, for example, ad  
7 market being the database where we store ads rather than  
8 user data, at least as of -- last time I touched this  
9 back in 2015, 2016, it would not be able to handle the  
10 load of storing user specific data. It's too small.  
11 But there are other larger systems that can handle user  
12 scale data such as Hive.

13 SPECIAL MASTER GARRIE: Yeah.

14 As I understand it, right, DIY has ad  
15 impressions, and that means every ad impression is  
16 associated somehow to the user, so that means it doesn't  
17 exist just in Hive.

18 MR. EUGENE ZARASHAW: I'm not sure how we  
19 produce the DIY file and whether that pulls from Hive  
20 or not.

21 SPECIAL MASTER GARRIE: Okay. Based on what  
22 I've heard from Facebook repeatedly -- but fair, good  
23 question, one to follow up on, and I will --

24 Steven, would you know by chance if the DIY  
25 file pulls from --

HIGHLY CONFIDENTIAL

Page 32

1 MR. STEVEN ELIA: I don't know that either.

2 SPECIAL MASTER GARRIE: Okay. Any next  
3 question is before we go down any other -- does Facebook  
4 have any legacy systems that contain plaintiffs -- that  
5 contain user data today? And I'll give you an example.

6 I used to work in banking for very large  
7 banks, like in building their technical systems. There  
8 were literally times where my clients would keep systems  
9 alive until the users would just die because nobody knew  
10 how the systems actually worked. So -- which is cheaper  
11 than trying to move these very wealthy individuals off  
12 these systems from a bank they acquired so they just  
13 kept them around. They called them legacy systems where  
14 people could access the data but not write new data to  
15 it.

16 So like does Facebook -- but I'm being  
17 specific, just broader. Does Facebook have any legacy  
18 systems that contain user data in it today?

19 Steven or Eugene, can you answer?

20 MR. EUGENE ZARASHAW: I think it might help to  
21 have a clear definition of legacy, Mr. Garrie, if that's  
22 possible. I don't think -- I don't know of any systems  
23 that would match the example you gave.

24 SPECIAL MASTER GARRIE: Oh, you don't do  
25 banking so I assume you have -- but I mean --

HIGHLY CONFIDENTIAL

Page 33

1 MR. EUGENE ZARASHAW: But the small number of  
2 users read only, don't write to system, and still being  
3 left up in production, I'm not familiar with a system  
4 quite like that.

5 SPECIAL MASTER GARRIE: Just products that  
6 sucked and you killed them; right? I mean, there's a  
7 lot of reasons why things may be legacy; right? Like --  
8 I mean, I think back to, you know, there had been lots  
9 of things, you know --

10 MR. EUGENE ZARASHAW: I think any point in  
11 time there are infrastructure systems we're actively  
12 trying to migrate off of.

13 SPECIAL MASTER GARRIE: Yeah.

14 MR. EUGENE ZARASHAW: But that would be in  
15 that list. It's usually about the process and they're  
16 still used to serve production traffic until they shut  
17 down.

18 SPECIAL MASTER GARRIE: Oh. Can you identify  
19 from that list which one those systems are? Because  
20 that would be very helpful.

21 MR. EUGENE ZARASHAW: I'm a bit out of date on  
22 what's currently being deprecated in infrastructure  
23 versus what's new and shiny.

24 SPECIAL MASTER GARRIE: Because, like, I mean  
25 I appreciate that you use elastic search but it's --

## HIGHLY CONFIDENTIAL

Page 34

1 that's a very -- people specialize their entire lives  
2 just in elastic search so it's worth noting.

3 So then it would be useful to know if Facebook  
4 has any legacy systems that have user data in them that  
5 are accessible through the DIY file.

6 MR. EUGENE ZARASHAW: I do not know the answer  
7 to that offhand.

8 SPECIAL MASTER GARRIE: Okay. And what I'm  
9 thinking about is remember Google Plus? I mean, I do;  
10 right? That was a horrible like -- well, it's all --  
11 Google Plus was a system Google made trying to do social  
12 media that failed epically.

13 And I'm willing to reckon that Google Plus --  
14 Google still has all of the Google Plus data, possibly.  
15 Maybe not. But that's an example of what I mean by a  
16 legacy system; right?

17 So Google doesn't run Google Plus today. But  
18 I would -- hypothetically they may still have all of  
19 that data. So that's what I'm referring to. Because  
20 there's been lots of machinations over the years of  
21 different attempts so...

22 MR. FALCONER: If that's not -- if that's not  
23 something that Mr. Zarashaw or Mr. Elia can answer  
24 confidently, that might be one where it makes sense for  
25 us to go back and make sure we've got the right answer.

HIGHLY CONFIDENTIAL

Page 35

1           SPECIAL MASTER GARRIE: That's fine. I just  
2 want to make sure we cover that. That's fine.

3           Mr. Elia, you're not in a position to answer  
4 this or are you?

5           MR. STEVEN ELIA: I'm not aware of others  
6 either.

7           I will add that as a general matter of policy,  
8 we don't retain data beyond the lifetime in which it's,  
9 you know, necessary and used.

10          So in these kind of scenarios where there's a  
11 product that got shut down and there's no way for a user  
12 to access it anymore, the expected behavior is that is  
13 not data we're going to keep around.

14          SPECIAL MASTER GARRIE: Yeah. But would you  
15 migrate -- is the data migrated to other systems then?

16          MR. STEVEN ELIA: If that product was, maybe,  
17 morphed into another one that would continue being  
18 accessible, that's possible. But if it was a product  
19 that users would have no way of getting back to, I  
20 wouldn't expect that data to be kept in any other form.

21          SPECIAL MASTER GARRIE: All right. So then  
22 I'd like to know if there are any of those legacy  
23 systems, just to confirm, and if there are, are they  
24 identified in the list? If they're not identified in  
25 the list, do they have user data? If they do have user

HIGHLY CONFIDENTIAL

Page 36

1 data, it would be helpful to explain what they do,  
2 rather than having to hold a whole nother hearing and  
3 back and forth, if possible.

4 So then I assume Facebook has a set of tools  
5 internal for customer sort and other things to track  
6 user data.

7 Are those tools identified in this list?

8 MR. EUGENE ZARASHAW: I don't know that we  
9 have tools to per se track user data.

10 SPECIAL MASTER GARRIE: You don't, like, have  
11 a customer support group that like gets access to --  
12 like take down my page, this isn't fair, they get a  
13 different view than -- of the user data than what the  
14 user sees?

15 MR. EUGENE ZARASHAW: We do have operations  
16 teams that manage specific queues of work that are task  
17 specific and they do get to see data about the user that  
18 is relevant to the specific query they're working.

19 So, for example, if there is -- the example  
20 you gave of a page being taken down, then the person  
21 working the appeal would have information about the  
22 page, some information about the user, the reason for  
23 the original decision that caused the page to be taken  
24 down and so on.

25 But it's not really for tracking other data

HIGHLY CONFIDENTIAL

Page 37

1 about the user. It would be the minimum data needed to  
2 make a decision on, in this case, whether to grant the  
3 appeal.

4 SPECIAL MASTER GARRIE: All right. Makes  
5 sense.

6 Are there other sort of systems like that?

7 MR. EUGENE ZARASHAW: There are a number of  
8 integrity systems that are similar to this where people  
9 losing access to their accounts, us taking down content  
10 for a variety of violations of policy, appeals for most  
11 of the take-down processes, similar systems for the  
12 developer platform, and so on.

13 SPECIAL MASTER GARRIE: And when you're  
14 looking at the user data, where is that user data coming  
15 from?

16 MR. EUGENE ZARASHAW: Generally from TAO.

17 SPECIAL MASTER GARRIE: Oh, from TAO. Okay.  
18 So we'll talk about TAO in a few minutes.

19 But where else does it come from besides --  
20 because TAO is an interface; right? So, I mean, it's a  
21 capability -- there's a lot -- you guys have your own  
22 have vocabulary here but we'll call it an interface  
23 capability, whatever you want to call it.

24 Besides TAO, where else is it getting from?

25 MR. EUGENE ZARASHAW: TAO is a primary storage

HIGHLY CONFIDENTIAL

Page 38

1 system. The majority of day-to-day use data on Facebook  
2 lives in TAO as its only storage.

3 TAO is underneath -- TAO is in itself a pretty  
4 complex storage application written on top of thousands  
5 of multiple databases but --

6 SPECIAL MASTER GARRIE: -- MySQL implement,  
7 right? You just abstracted it out.

8 MR. EUGENE ZARASHAW: Exactly. It's a grab  
9 database at the top of MySQL.

10 SPECIAL MASTER GARRIE: Okay. Now we can --  
11 technical terms are better for me than nontechnical.

12 So okay.

13 MR. EUGENE ZARASHAW: With replication and all  
14 that built -- handled properly across data centers so  
15 multiple things.

16 SPECIAL MASTER GARRIE: Okay. So TAO is  
17 backed by MySQL. So that's where all the user data is  
18 stored.

19 MR. EUGENE ZARASHAW: I may be out of date.  
20 I think we may have migrated the backing store for TAO  
21 to a different database. I'm out of date on this one.

22 Maybe something faster than MySQL.

23 But, yes, most user data, including data about  
24 a page, why we took down the page, the user who is --  
25 had been on the page, the primary storage for that would



## HIGHLY CONFIDENTIAL

Page 39

1 be TAO.

2 SPECIAL MASTER GARRIE: Okay. So then does  
3 Facebook store data about users that is not directly  
4 associated with the user or -- I don't know.

5 What do you guys call it? What's the key for  
6 a user? You called it a token.

7 Is that -- are users -- so there's a lot of  
8 terms that have been used so can we agree on one?

9 Facebook ID? Token? Whatever? The key?

10 MR. EUGENE ZARASHAW: The primary key would be  
11 a user ID.

12 SPECIAL MASTER GARRIE: Okay. So then my  
13 question is does Facebook store data about users not  
14 directly associated with a Facebook ID but is capable  
15 of being associated with that user?

16 MR. EUGENE ZARASHAW: I -- probably, in that  
17 you -- we store data that is based on IP addresses, and  
18 with sufficient work, you could be -- you could make a  
19 probable guess between a timestamp and an IP address,  
20 what user was using that IP address, not -- not with  
21 certainty but with probability.

22 SPECIAL MASTER GARRIE: Yeah. I mean, the  
23 only thing certain is I hold this pen and I drop it, you  
24 know it hits the ground; right?

25 So -- so then I guess address, e-mail, phone

HIGHLY CONFIDENTIAL

Page 40

1 number. Okay. That makes sense.

2 So then what is the process that enables data  
3 that is not stored in a manner directly -- so what  
4 systems -- what's the process or systems where -- that  
5 enable data that's not stored in a manner directly  
6 associated with the Facebook ID to be associated with  
7 it?

8 Is that like what you're using -- sorry.

9 MR. EUGENE ZARASHAW: I don't know of any  
10 single process because it -- there's a multitude of  
11 systems that store data in different ways, some of which  
12 don't store it keyed off a user ID, in which case -- if  
13 a system is not designed to pull data by user ID, it  
14 would take a significant amount of work to potentially  
15 modify the system to be able to look it up by user ID.

16 The analogy -- oh, sorry.

17 SPECIAL MASTER GARRIE: No. Go ahead.

18 MR. EUGENE ZARASHAW: The analogy I would use  
19 is suppose you have an archive of newspapers that are  
20 indexed by the language of the newspaper and the date of  
21 that edition, and you want to find out whether my name  
22 was mentioned in it.

23 It would require going through the entirety of  
24 the archive quite manually or writing new software to go  
25 through the archive to find mentions of my name because

HIGHLY CONFIDENTIAL

Page 41

1 during the original design, it was designed to only look  
2 things up by date and language.

3 So we have a number of systems like that where  
4 it was just not designed from day one to serve queries  
5 of that type.

6 SPECIAL MASTER GARRIE: So then -- sorry. I'm  
7 just thinking about it all. Sorry. I'm listening.

8 So -- but you have a -- using your analogy,  
9 you have a list of delivered papers with addresses;  
10 right?

11 MR. EUGENE ZARASHAW: Often not even  
12 addresses. Often something much more rudimentary like  
13 some other identifier --

14 SPECIAL MASTER GARRIE: You could delivery the  
15 paper if it rained that day.

16 But the point is you have a list of delivered  
17 papers with some identifier, the address or whatever,  
18 and then you connect the paper to a -- you can connect  
19 that paper to a user.

20 MR. EUGENE ZARASHAW: Sometimes. Sometimes we  
21 don't even maintain the identifiers. A lot of it ends  
22 up being anonymized intentionally or unintentionally  
23 because only the necessary data is maintained.

24 So, for example, you might store performance  
25 data about a system with no useful identifiers at all.

HIGHLY CONFIDENTIAL

Page 42

1 Just how did the system perform during each request with  
2 no IP address and no user data, which would make it  
3 nearly impossible to tie it back.

4 SPECIAL MASTER GARRIE: Okay. So then if we  
5 look at the systems, these 43 other systems, so if  
6 you select -- you select -- you remove the 12, do you  
7 have any idea of what of those 43 would be -- right? I  
8 mean, they all have data. The question is what -- how  
9 is that data being, you know, used -- or connected to --  
10 using your newspaper example, how do we know which ones  
11 of those systems actually can or use to take the list of  
12 delivered papers with addresses and connect the papers  
13 to a user, as an example?

14 MR. EUGENE ZARASHAW: It would take a very  
15 significant number of engineers a fair bit of time to  
16 even answer that question, unfortunately. I do not know  
17 offhand.

18 SPECIAL MASTER GARRIE: But then the engineers  
19 for each of those systems, I mean, at some level someone  
20 has to know -- I mean, let me be more tangible.

21 If we look at a system like Dumbo -- I mean, I  
22 don't know what Dumbo is. It's literally System No. 6,  
23 Dumbo --

24 MR. EUGENE ZARASHAW: I don't either.

25 SPECIAL MASTER GARRIE: Okay. Steven, do you

HIGHLY CONFIDENTIAL

Page 43

1 know what Dumbo is?

2 MR. STEVEN ELIA: I don't, and I don't even  
3 recognize the majority of the list.

4 SPECIAL MASTER GARRIE: Yeah.

5 Oh. What about -- oh, payments is a MySQL  
6 database.

7 Is there a user data in payment -- I assume,  
8 taking our newspaper address example, and payments is  
9 No. 38 in the list -- well, I'm saying list. There's a  
10 letter provided by Facebook called Facebook Supplemental  
11 Submission. I could share it. I mean, it -- but  
12 there's this list of systems identified.

13 Would you know if payments is one of those  
14 systems we're talking about?

15 Is there any way if you look at this, you and  
16 Steven -- I can share my screen -- say which ones would  
17 possibly be able to do that newspaper example?

18 MR. STEVEN ELIA: I think it would be  
19 difficult to say which ones might be, given that Eugene  
20 and I don't recognize many of these.

21 The payments example I believe was one of the  
22 12 later identified, so that may not be the best example  
23 to ask us on.

24 SPECIAL MASTER GARRIE: Oh, yeah. Sorry. I  
25 tried to highlight them all.

HIGHLY CONFIDENTIAL

Page 44

1                   What about Lldbpostmortem Cache? It's No. 48.

2                   MR. EUGENE ZARASHAW: I have not heard of this  
3 one.

4                   MR. STEVEN ELIA: Same.

5                   SPECIAL MASTER GARRIE: How about  
6 Research Home; is that one? I mean, that... What's  
7 Research Home?

8                   Is that one you can connect back?

9                   MR. EUGENE ZARASHAW: I don't know of that one  
10 either. I mean, I can volunteer an example, if you'd  
11 like, that might be a good one, which is --

12                  SPECIAL MASTER GARRIE: Yeah.

13                  MR. EUGENE ZARASHAW: -- Unicorn.

14                  SPECIAL MASTER GARRIE: Unicorn. Okay.  
15 Didn't think Unicorn. All right. Good.

16                  MR. EUGENE ZARASHAW: So Unicorn is a search  
17 back end for searching. And it's a duplication of data  
18 stored in other systems. So it enables the -- using  
19 Unicorn to index another system allows you to create a  
20 search application for another system -- for specific  
21 types of queries for another system. So it's generally  
22 pointed at things inside TAO, because TAO itself does  
23 not serve any kind of searches.

24                  So there could possibly -- depending on what  
25 kind of search indexes somebody at one point created, it

HIGHLY CONFIDENTIAL

Page 45

1 is possible for there to be user data inside Unicorn,  
2 but it would be duplicated in the other systems that  
3 Unicorn is fronting.

4 SPECIAL MASTER GARRIE: Because it's sitting  
5 over.

6 So then if they -- so then if you infer data  
7 about the users and it could be connected to people --  
8 all right. That makes sense.

9 So then what I'm trying to figure out is one  
10 of these systems -- so then you have -- because then, as  
11 I understand it, right -- I guess maybe the better way  
12 to look at it is what systems store the information used  
13 to create each Facebook ID?

14 Can we do it that way?

15 Like a user profile?

16 MR. EUGENE ZARASHAW: What do you mean create?

17 SPECIAL MASTER GARRIE: Well, when you create  
18 a new user, right, it stores data?

19 MR. EUGENE ZARASHAW: TAO.

20 SPECIAL MASTER GARRIE: Where is that -- on  
21 TAO.

22 Because -- so then -- so then what systems  
23 store the data necessary from that list to compile a  
24 single user's profile or profile once it's created?

25 MR. EUGENE ZARASHAW: By compile, do you mean

HIGHLY CONFIDENTIAL

Page 46

1 for DIY or for a certain type --

2 SPECIAL MASTER GARRIE: Not DIY. I'm talking  
3 for Facebook. Because remember the definition the  
4 Court -- so, yes, I wish that was that simple.

5 The Court said data collected from a user's  
6 on-platform activity, data obtained from third parties  
7 regarding a user's off-platform activities, and data  
8 inferred from a user's on-or-off platform activity. So  
9 it's much broader than what sits in the DIY file.

10 MR. EUGENE ZARASHAW: I don't have a good  
11 answer to that one as to which of these systems may have  
12 some of it.

13 The hard part of that question is also knowing  
14 which parts are just duplicates of something else.

15 SPECIAL MASTER GARRIE: So if I could get the  
16 first question answered, then I want the second part  
17 answered.

18 But it's very challenging -- who at -- if it's  
19 not you or Steven that can answer this, who at Facebook  
20 can -- 'cause it's pretty simple; right? I mean, this  
21 question is simple but the answer is deceptively  
22 complicated. And I'm just trying to understand at the  
23 most basic level from this list what we're looking at.

24 MR. EUGENE ZARASHAW: I don't believe there's  
25 a single person that exists who could answer that



## HIGHLY CONFIDENTIAL

Page 47

1 question. It would take a significant team effort to  
2 even be able to answer that question.

3 SPECIAL MASTER GARRIE: Okay. There's a  
4 Facebook ID and there's an ad impression mapping; right?  
5 'Cause we know that because it's in the DIY file; right?

6 It's either in the Hive or it's extracted  
7 periodically and stored elsewhere; right?

8 MR. EUGENE ZARASHAW: (Nodding.)

9 SPECIAL MASTER GARRIE: Where is it stored --  
10 who knows where these things are being stored?

11 MR. EUGENE ZARASHAW: Whoever is responsible  
12 in the ads team for logging impressions would be the  
13 most likely person to know everywhere it would be  
14 logged.

15 SPECIAL MASTER GARRIE: And that would be  
16 associated with a user?

17 MR. EUGENE ZARASHAW: And which parts of the  
18 logging are user associated versus which parts of the  
19 logging are only associated with the ad in advertiser  
20 but not the user.

21 SPECIAL MASTER GARRIE: Yeah.

22 So --

23 MR. EUGENE ZARASHAW: And the challenge is for  
24 each product, it would be a different set of people that  
25 would know how that specific product area does logging

HIGHLY CONFIDENTIAL

Page 48

1 and data retention.

2 SPECIAL MASTER GARRIE: Because I see these  
3 systems that have ad in it but nobody seems to be able  
4 to tell me if it's user data in the ad system.

5 And I know that you can infer and make  
6 relationships about the user ad data that's -- because  
7 you have it in your DIY file, but there's obviously more  
8 data that gets attached to the user's activity to the  
9 ad unit that's sitting somewhere in these systems. And  
10 I'd really like to not make Facebook go and figure it  
11 out for every single system. But the Court says  
12 on-platform or off-platform activity, so associated with  
13 a user.

14 So -- but I hear you.

15 All right. Steven, do you have any idea?

16 MR. STEVEN ELIA: As I said before, I also  
17 don't even recognize some of these. I would also agree  
18 there's not a single individual who would recognize all  
19 of these or be familiar enough with all of these.

20 SPECIAL MASTER GARRIE: But there are -- there  
21 would be a group of individuals that could access --  
22 someone at some level -- because the ad unit is a great  
23 example; right?

24 Facebook ID to ad impression maps; right?

25 It's either in the Hive and extracted in a periodic

HIGHLY CONFIDENTIAL

Page 49

1 fashion and stored elsewhere because that's how it ends  
2 up in the DIY file.

3 Whatever system it's extracting it from would  
4 theoretically have a richer dataset about the ad user  
5 engagement, or maybe not, but someone needs to clarify  
6 of your list of systems yes or no.

7 MR. STEVEN ELIA: What I would say is each of  
8 these 55 systems will have a team of engineers who are  
9 responsible and knowledgeable about it, but even then,  
10 those individuals may not know all of the ways in which  
11 that system is used by other teams.

12 SPECIAL MASTER GARRIE: I get that.

13 But I'm looking at it from the -- yes. I  
14 agree. But I'm looking at it in the narrow definition  
15 of how the Court has defined user data as to what needs  
16 to be pulled for these users.

17 And the Court has made a very explicit  
18 determination, Judge Corley and Judge Chhabria  
19 affirmed -- I don't know who did. But Judge Corley  
20 clearly said these three things; right?

21 And so there has to be some way to -- like I  
22 have no problem with an -- like if it's too hard or too  
23 complicated, you could articulate them. But what it  
24 can't be is it requires a group of engineers to tell me  
25 whether or not the ad unit and the ad impression mapping

HIGHLY CONFIDENTIAL

Page 50

1 part of it, you know, because it's going take a long  
2 time and cost a lot of money. Like I appreciate all  
3 that but there needs to be tangible pieces of it.

4 But if you don't know the answer, that's a  
5 legal argument which has no business here. This is just  
6 more just I want to get as many technical questions  
7 cleared up. And I can bring my legal questions up with  
8 the parties downstream and not waste your time.

9 MR. FALCONER: I think what -- sorry. I don't  
10 mean to cut you off.

11 But for Mr. Garrie, I don't think Mr. Zarashaw  
12 or Mr. Elia works a lot with the ads -- business in the  
13 ads platform.

14 So maybe -- could we take that one back, too,  
15 of -- to find someone who works in this area a little  
16 bit more frequently and directly and see if that might  
17 be one where we could get back to you with a better  
18 answer from someone who is more familiar with this in  
19 their day-to-day work?

20 SPECIAL MASTER GARRIE: Because all I'm trying  
21 to do is go through this list and figure out what  
22 systems have user data and how they -- and how --  
23 like -- because I'm pretty certain that Facebook --  
24 the -- Facebook ID, just based on my understanding of  
25 the DIY file, is pulling data from different systems.

## HIGHLY CONFIDENTIAL

Page 51

1 And those systems may have additional data about the  
2 user sitting in them. They may or may not be  
3 accessible, or may be duplicative or may not be, but  
4 there has to be some way of figuring out the  
5 information; right?

6 Because IP addresses may sit in this  
7 database -- or ad impressions -- I don't know about ad  
8 impressions -- but my likes sit in this database, and  
9 the DIY file may -- the application may pull from that.  
10 But the database may have other information about my --  
11 the user.

12 And the Court, just with this definition,  
13 intentional or otherwise, is on platform activity, I  
14 mean, it's -- I clicked a like; right? I mean, it is --  
15 as an example. But we'll keep going.

16 So then does Facebook include information in a  
17 user profile that is not posted by the user to Facebook?

18 MR. EUGENE ZARASHAW: What definition should  
19 we use for the user profile in this case?

20 SPECIAL MASTER GARRIE: Well, I could tell you  
21 how the Court -- well, interesting you ask because you  
22 have patents that define how you create user profiles  
23 that I've actually read.

24 But we won't get into the patent piece of it  
25 and just --

HIGHLY CONFIDENTIAL

Page 52

1                   How would you define a user profile?

2           A Facebook ID; right? That's a human -- how about this,  
3           a breathing human entity interacting with your system?

4                   MR. EUGENE ZARASHAW: But what I mean is what  
5           else makes up a user profile?

6                   We don't generally -- when we look at user  
7           profile internally, that term generally means a user,  
8           but it's not necessarily a collection of data. So I'm  
9           not sure how to answer that question.

10                   SPECIAL MASTER GARRIE: So the user is a -- I  
11          mean, the user, I guess you want to call it an object.  
12          I mean, calling it -- I mean, the user is -- you have a  
13          whole social graph schema. Like, I mean, you clearly --  
14          there's a user -- I'm not as familiar as you're -- as  
15          you guys are of your system, but we'll call the user an  
16          object that interacts with a whole litany of services  
17          and systems.

18                   Some data, right, like some is  
19          system-collected data, like the mobile activity, where  
20          they were when they logged in, their OS, whatever, or  
21          the computer or whatever.

22                   But what you -- I guess the best way to think  
23          about it is if you say Facebook include -- and its user  
24          profile that is not posted by the user to Facebook so  
25          the user -- I guess the Facebook profile page, if you

HIGHLY CONFIDENTIAL

Page 53

1 want to -- we can be -- maybe that helps -- let's say  
2 Facebook profile page. Let's start with that.

3 MR. EUGENE ZARASHAW: So is the question then  
4 is anything rendered on the Facebook profile page for a  
5 user that was not put there by the user?

6 If so, yes, in that --

7 SPECIAL MASTER GARRIE: Yes.

8 MR. EUGENE ZARASHAW: -- there would also be  
9 activity of other users that involve this user's  
10 activity.

11 So as an example, if user A tags user B in a  
12 photo, that photo may appear on user B's page, even  
13 though user B did not post it.

14 SPECIAL MASTER GARRIE: So then how is -- so  
15 is there a list of what information that is and how it's  
16 obtained?

17 MR. EUGENE ZARASHAW: The challenge with  
18 answering that question is the data is not -- you can't  
19 always go from a user to all data referencing a user.

20 SPECIAL MASTER GARRIE: No. We talked about  
21 that. I get that.

22 Let's just say Facebook profile page is how  
23 we'll start with.

24 So does Facebook include information in  
25 Facebook profile page that is not posted by the user in

HIGHLY CONFIDENTIAL

Page 54

1 Facebook? The answer is yes.

2 And then my question is, do you have a list of  
3 what information that is?

4 MR. EUGENE ZARASHAW: I do not have a complete  
5 list, but some of that would be in TAO and some of that  
6 would be in TimelineDB for the case of the user profile.

7 And then I don't recall which system we use  
8 these days to store the photos, but if the photo -- if a  
9 user was tagged in a photo, the photo itself would be  
10 stored in yet some other database. And then the  
11 identifier for the photo stored in that database would  
12 be stored with the post inside the timeline most likely.

13 And speaking of the earlier challenge, you --  
14 I don't believe you would be able to go from the photo  
15 to the user. Only the other direction -- or from the  
16 photo to the post, but only in the reverse direction.

17 SPECIAL MASTER GARRIE: Okay. That makes  
18 sense.

19 So then --

20 MR. FALCONER: Mr. Garrie, I don't know if  
21 you're about to change topics. We've been going --

22 SPECIAL MASTER GARRIE: I am.

23 MR. FALCONER: -- a bit more than an hour.  
24 Would it be okay if we took --

25 SPECIAL MASTER GARRIE: Yeah. We can take a



## HIGHLY CONFIDENTIAL

Page 55

1 break. Sorry. I get --

2 MR. FALCONER: -- so Mr. Elia and Mr. Zarashaw  
3 can stretch their legs, get some water, and use the  
4 bathroom.

5 SPECIAL MASTER GARRIE: Yeah. I apologize.  
6 Certainly. So literally it's a perfect breaking point.  
7 We're -- yeah. We can take a five-minute break.

8 Plaintiffs, if you want to -- if you have any  
9 questions that you want to ask, I'm happy to hear them  
10 when we come back from our five-minute break. We'll go  
11 off -- wait, before we go off the record. Sorry. We'll  
12 go -- I'm sorry.

13 So the plaintiffs, organize your questions  
14 during the five-minute break. We'll go off the record  
15 for five minutes and we'll resume -- we're actually  
16 right at the breaking point so that's perfect.

17 And then if Facebook -- I'm sorry. If counsel  
18 for Facebook has any follow-up questions that they think  
19 should be asked, please feel free to ask -- you know,  
20 spend your five minutes thinking of any of those  
21 questions as well. And we'll resume in five minutes.

22 Okay?

23 So now we're going to go off the record.

24 (Break taken in proceedings.)

25 SPECIAL MASTER GARRIE: So plaintiffs, are

HIGHLY CONFIDENTIAL

Page 56

1     there any questions that you want to ask me that you  
2     think I should ask or follow up based on what we've  
3     discussed so far?

4             MS. WEAVER: Yes. I think there's one general  
5     framing question and then there's a follow-up with some  
6     specific questions about identifiers.

7             The first general framing is trying to find  
8     the internal word or jargon that Facebook uses for what  
9     externally people refer to as a profile. And to  
10    distinguish that from the profile that is on, you know,  
11    a person's Facebook page, which is to say:

12            When Facebook is seeking to provide  
13    personalized ads to Facebook's actual customers by  
14    drawing off datasets about, say, one of our named  
15    plaintiffs, they aren't going to the download your  
16    information tool. They are drawing from some different  
17    kinds of datasets that is contained within Facebook that  
18    is much larger in scope than just what somebody posted.

19            It includes that inferred data from on and off  
20    the web that's in one of the three categories you named.

21            And so the question is -- that's the framing.  
22    The specific question is: Are there identifiers or  
23    cookies, like the data cookie or a pixel, that can be  
24    used to trace data connected to the named plaintiffs in  
25    realtime today that we could use as an exemplar?

HIGHLY CONFIDENTIAL

Page 57

1           For example, if we had today the named  
2       plaintiff go on, we looked at the data flow and said  
3       there's a Facebook data cookie that is associated with  
4       them, could we go in realtime today and use an  
5       identifier like that to pull information that is  
6       inferred? Like the -- whatever verticals or segments  
7       these named plaintiffs are in.

8           Not the generic ones that are available on the  
9       outward-facing Facebook profile, but the internal ones  
10      that Facebook uses to personalize targeted ads.

11           SPECIAL MASTER GARRIE: That was a topic of my  
12      next section, but we can kick it off.

13           Could you reask the question, though --  
14           Go ahead, Counsel Falconer. Did you have any  
15      question?

16           MR. FALCONER: Well, just -- I'm not sure if  
17      Mr. Elia is back yet.

18           SPECIAL MASTER GARRIE: Oh, he's not. I'm  
19      not --

20           So, Carlos, could you find Mr. Elia, or  
21      Steven -- or Mr. Elia and move him back?

22           MR. FALCONER: I also --

23           JAMS CONCIERGE CARLOS ESPINOZA: I think he  
24      might have left the session. I don't see him listed  
25      anywhere.

HIGHLY CONFIDENTIAL

Page 58

1           SPECIAL MASTER GARRIE: That could be a result  
2 of me, to be honest.

3           MR. FALCONER: I also am not on video for  
4 everyone so I need -- I apologize. I need to log out  
5 and log back in because I can't see anybody's video. So  
6 if that's all right, I'll do that as quick as I can.

7           SPECIAL MASTER GARRIE: Okay.

8           MR. SANDEEP SOLANKI: I can try to ping Steven  
9 just on our internal chat. So I assume that's okay with  
10 everybody, just to see where he is.

11          SPECIAL MASTER GARRIE: Yes. We can go off  
12 the record.

13          MS. WEAVER: Special Master Garrie, we know  
14 that there's a replacement ID. You know, is that an  
15 identifier that can be used to track people? Is it the  
16 data cookie? I don't really know. Are there cookie  
17 values we could create right now that are connected to a  
18 Facebook user ID to use as an exemplar?

19          SPECIAL MASTER GARRIE: So I got the question  
20 gist.

21                So it is the next section that I sort of have  
22 to cover with them. So I will get to your question as  
23 it comes up in a timely manner as we go into this next  
24 sort of section.

25                But I have similar -- so -- questions there,

HIGHLY CONFIDENTIAL

Page 59

1 so...

2 Why don't we give them a minute to...

3 MR. SANDEEP SOLANKI: Oh, he said the Zoom  
4 disconnected, gives him an error to log back into it.

5 (Pause in proceedings.)

6 SPECIAL MASTER GARRIE: We can go back on the  
7 record.

8 And I can ask Mr. Zarashaw my first question  
9 in this area and when -- I'm assuming if his answer is  
10 very succinct, we'll go back off the record. If not, we  
11 will continue forward. And hopefully Mr. Elia will be  
12 back.

13 With that in mind, going back on the record, I  
14 have a general question, which is sort of how does  
15 Facebook refer to the set of data it has about users  
16 inclusive of embedded and inferred data associated with  
17 a particular individual?

18 MR. EUGENE ZARASHAW: The only term I could  
19 think of is user data.

20 Just to add to that, we don't generally talk  
21 about it in terms of -- in those terms of user and all  
22 the data associated with them. Most of our  
23 conversations are in the other direction of a product  
24 and the things related to making the product -- specific  
25 product work.

## HIGHLY CONFIDENTIAL

Page 60

1           SPECIAL MASTER GARRIE: As engineers, when you  
2 build something, how do you refer to it?

3           MR. EUGENE ZARASHAW: Honestly, data.

4           SPECIAL MASTER GARRIE: Yeah. But I'm a  
5 unique piece of data. You call me a user object.

6           How do you -- there has to be some mechanism  
7 for when you're developing things to refer to the data  
8 about a user.

9           MR. EUGENE ZARASHAW: So we generally refer to  
10 them as users or as people, and a user ID is the primary  
11 identifier for a user.

12           I think the challenge is in the direction of  
13 the connections between the data. A lot of -- in my  
14 experience outside of Facebook, a lot of the time it --  
15 things would start with a user and fan out from the  
16 user. So you could start with a user -- a simple  
17 product and find all the data about them.

18           Because we have centralized infrastructure for  
19 certain things, and then many products and surfaces  
20 built that both share some of the same infrastructure,  
21 it tends to be the other direction around where you have  
22 additional data in various systems that reference a  
23 user, but it's -- unless you know what you're looking  
24 for --

25           As an example, it's much easier if you have a

## HIGHLY CONFIDENTIAL

Page 61

1 post to go find the user that made the post than if you  
2 have a user to find all of the posts they made.

3 SPECIAL MASTER GARRIE: Like a phone number or  
4 something key to the user?

5 MR. EUGENE ZARASHAW: The key is the user ID.  
6 It's a 64-bit integer that is the unique identifier for  
7 the user that is used for most systems inside Facebook  
8 that are not anonymized.

9 SPECIAL MASTER GARRIE: That's more than all  
10 the people in the world so I'm pretty sure you don't run  
11 out of numbers.

12 MR. EUGENE ZARASHAW: Funny topic, we actually  
13 are.

14 SPECIAL MASTER GARRIE: Really?

15 MR. EUGENE ZARASHAW: Because we share that  
16 64-bit ID space with other objects as well, due to some  
17 poor infrastructure choices a long time ago.

18 So we're running out of 64-bit numbers.

19 SPECIAL MASTER GARRIE: Got it.

20 So then the 64-bit numbers, every user that  
21 trace through -- okay.

22 MR. EUGENE ZARASHAW: But there are systems  
23 that don't care about the user in question, so the user  
24 ID can get dropped often as the data traverses through  
25 systems.

HIGHLY CONFIDENTIAL

Page 62

1 SPECIAL MASTER GARRIE: Okay.

2 MR. EUGENE ZARASHAW: Basically the general  
3 privacy take on this is if we don't need to log PII, we  
4 don't.

5 SPECIAL MASTER GARRIE: No. I get that.

6 But how then -- there has to be some mechanism  
7 that I could use to search these different systems to --  
8 because I'm at a -- like, the Court's definition, right,  
9 unfortunately says "Data collected from a user's  
10 on-platform activity," right, which could -- and the  
11 second is "Data obtained from third parties regarding a  
12 user's off-platform activities," which we're about to  
13 get to, and third is "Data inferred from a user's  
14 on-or-off platform activities."

15 So the issue then becomes all of these  
16 systems, there has to be some -- a mechanism that will  
17 allow me to trace the data that is associated with that  
18 64 bit key.

19 MR. EUGENE ZARASHAW: I think that's the  
20 challenge we're facing right now is that is not a use --  
21 most of these systems exist because the generalized  
22 systems weren't sufficient usually for performance  
23 reasons.

24 So the reason you're seeing 55 systems on that  
25 list is because we started out with two or three, and



HIGHLY CONFIDENTIAL

Page 63

1 for various reasons, as we built certain products, the  
2 performance or cost characteristics of the generalized  
3 systems just didn't work for those use cases.

4 And so we ended up adding specific systems to  
5 serve just one use case.

6 One example, would be storage of large blobs  
7 of data, i.e., photos or videos, is highly specialized  
8 and only does one thing. It does it far better than our  
9 other methods of storage but it only does that one  
10 thing.

11 The use case you're mentioning of I would like  
12 to look up all of a thing about a user from the system  
13 just was not in this original specification, and it  
14 would have to be built.

15 So I actually think the photo system is a  
16 great example because it would be a rather nontrivial  
17 effort if you went into that system and tried to find  
18 old photos for a user.

19 You would have to go to a completely different  
20 system, find how the photos are tracked about a user,  
21 get the entire list of photo identifiers, then go back  
22 to the system and look up each photo.

23 Just because we never originally built the  
24 capability to go and find all photos by a user key. In  
25 fact, I'm not certain a user key exists in that system

HIGHLY CONFIDENTIAL

Page 64

1 at all.

2 I think most of the challenges we're going to  
3 see is because of these highly specialized systems where  
4 it was a implementation cost and operational cost  
5 decision to only build for the specific use case they  
6 are good at and nothing else.

7 SPECIAL MASTER GARRIE: Then, Steven, I'm  
8 going to ask you the same question.

9 MR. STEVEN ELIA: Sorry. I had technical  
10 issues. Could you repeat the --

11 SPECIAL MASTER GARRIE: I'll reask the  
12 question. Yes, of course.

13 So how does Facebook refer to the set of data  
14 it has about users, inclusive of embedded and inferred  
15 data associated with the individual user? Like a  
16 Facebook profile; right? I mean, how does Facebook  
17 refer to it?

18 MR. STEVEN ELIA: I would refer to a profile,  
19 a profile data in a few ways myself, and that's probably  
20 in a much more narrow fashion than I think the Court had  
21 defined here.

22 Often when I'm thinking about a profile, I  
23 think of the data directly on that user object in TAO  
24 and other data like their content that is associated  
25 with that, usually in that same system.

HIGHLY CONFIDENTIAL

Page 65

1           SPECIAL MASTER GARRIE: Then what about the  
2       inferred data that you get associated with that user ID?

3           MR. STEVEN ELIA: Do you have an example you  
4       have in mind of inferred data?

5           SPECIAL MASTER GARRIE: I mean, you collect  
6       on -- there's a litany of data; gender, age. I mean, we  
7       can -- I can give a very -- I guess maybe the newspaper  
8       example, you know, where we have a list of the delivered  
9       papers with addresses, and I can connect that paper to  
10      that user. And then from there, I could -- you know, or  
11      ad impressions; right? You know...

12           We could use the newspaper. Maybe that's a  
13      good place to start.

14           MR. STEVEN ELIA: As Eugene was saying, I  
15      think -- I only caught part of his answer there before I  
16      reconnected -- when we don't store the user ID with some  
17      of this data, it's generally very difficult to look that  
18      up by user or associate with the user because that is  
19      not indexed in that fashion.

20           And so when I, again, think of this kind of  
21      user data, I would think of, well, what is the data  
22      that's primarily indexed by that user ID?

23           SPECIAL MASTER GARRIE: You have 55 systems.

24           MR. STEVEN ELIA: Many of which I wouldn't  
25      expect to contain user IDs at all or even if they did,

HIGHLY CONFIDENTIAL

Page 66

1       certainly not be indexed by that.

2               SPECIAL MASTER GARRIE:   Well -- go ahead.

3               MR. STEVEN ELIA:   Yeah.   And so connecting any  
4       data in there back to a particular user would be  
5       difficult or impossible in most cases, if that's not how  
6       it's stored.

7               SPECIAL MASTER GARRIE:   But a system that  
8       takes a user ID, right, and then returns data that  
9       matches to the user somehow is based on characteristics.

10              MR. STEVEN ELIA:   A system that takes in a  
11       user ID and indexes the data that way, those systems do  
12       exist, but most of the ones in that list of 55, I don't  
13       believe operate in that kind of fashion.

14              SPECIAL MASTER GARRIE:   Can you identify which  
15       ones you know of on that list that operate in that  
16       fashion?

17              MR. STEVEN ELIA:   The primary system that  
18       operates that way is TAO, and that's the system -- the  
19       only one I can think of that I've personally used that  
20       operates in that fashion.

21              MR. EUGENE ZARASHAW:   The other one I can  
22       think of would be Timeline.

23              SPECIAL MASTER GARRIE:   MySQL doesn't work  
24       with the user ID key?

25              MR. EUGENE ZARASHAW:   Depends how the keys are

HIGHLY CONFIDENTIAL

Page 67

1 set up. TAO is a layer on top of it that does its own  
2 keys.

3 SPECIAL MASTER GARRIE: That's' a question I  
4 had.

5 Why is MySQL in there?

6 MR. EUGENE ZARASHAW: We still have some MySQL  
7 servers that are not back in TAO. They are just used  
8 directly. They are also generally not used to use store  
9 user scale data.

10 SPECIAL MASTER GARRIE: Does payments use a --  
11 like some of these I would think have some user ID, like  
12 payments or order DB or ads online storage.

13 MR. EUGENE ZARASHAW: There are some  
14 databases, Clusters -- such as payments, that are  
15 clusters of MySQL that are not based in TAO.

16 I think one thing that may help, given you  
17 worked in banking tech, is TAO doesn't have  
18 transactions. So any time you want transaction support  
19 such as payments, TAO becomes problematic.

20 SPECIAL MASTER GARRIE: I lost connectivity  
21 here.

22 How is that possible?

23 Is this the internet? Oh, you gotta be  
24 kidding me.

25 Can you hear me, Kathryn?

HIGHLY CONFIDENTIAL

Page 68

1 THE COURT REPORTER: I can hear you, yes.

2 SPECIAL MASTER GARRIE: Okay. It skipped out.

3 Can you send me the link for the realtime  
4 transcript so I can read it as well?

5 THE COURT REPORTER: Sure. Can I have your  
6 e-mail?

7 SPECIAL MASTER GARRIE: Yeah. D -- I'll put  
8 it in chat for you.

9 So what are those clusters, not the TAO ones,  
10 I guess then?

11 MR. EUGENE ZARASHAW: TAO -- in some cases  
12 because [REDACTED]

13 [REDACTED]. So I think you dropped  
14 out during -- uh-oh.

15 You dropped out during the last answer. But I  
16 mentioned payments, for example, as per your banking  
17 days, you had -- you had databases that had transaction  
18 support.

19 [REDACTED]

20 [REDACTED]

21 [REDACTED]

22 SPECIAL MASTER GARRIE: I mean, I guess  
23 plaintiff's question -- so you're saying there's no  
24 unifying user ID or key that goes across all these  
25 systems?

HIGHLY CONFIDENTIAL

Page 69

1 MR. EUGENE ZARASHAW: Across all of them, no.

2 The most common way to represent a user is a  
3 user ID. However, the user ID may not be associated  
4 with a lot of the data stored in many IT systems.

5 SPECIAL MASTER GARRIE: Okay. All right.  
6 We'll return to this once I've processed it a little  
7 further.

8 Off platform -- so my next question is: Does  
9 Facebook -- and I don't know if this is for Mr. Elia or  
10 you Mr. Zarashaw -- but does Facebook receive data about  
11 user's off-platform activity?

12 MR. EUGENE ZARASHAW: We receive some data  
13 that I know of due to us having the like button, and the  
14 Facebook Javascript SDK. So some data does come in  
15 about the source of where this was loaded up or somebody  
16 liking a page off of Facebook and so on.

17 SPECIAL MASTER GARRIE: I took a look at the  
18 SDK. It's a little broader than that for your inbound.

19 Do you have a list of data you do receive  
20 about users and then -- so I can be specific. User  
21 tracking cookies?

22 MR. EUGENE ZARASHAW: I do not know the  
23 details of the ad side of the SDK. The last time I  
24 interacted with the SDK would have been about 2013,  
25 maybe 2014.

HIGHLY CONFIDENTIAL

Page 70

1 SPECIAL MASTER GARRIE: Got it.

2 What about you, Steven? I can list them off  
3 if you guys -- if either one of you know the answers.

4 User tracking cookies?

5 Mr. Elia?

6 MR. STEVEN ELIA: I don't know. I'm not  
7 familiar with the use of cookies. No.

8 SPECIAL MASTER GARRIE: All right. What about  
9 Facebook pixel?

10 MR. STEVEN ELIA: I believe that's for ads  
11 purposes, which I haven't worked on either.

12 SPECIAL MASTER GARRIE: Mr. Zarashaw?

13 MR. EUGENE ZARASHAW: Same issue. I'm very  
14 loosely familiar with the pixel and not the SDK details.

15 SPECIAL MASTER GARRIE: Is it possible that  
16 the Facebook pixel, plus the user ID would give us a  
17 decent idea of what systems have what user ID -- data?  
18 Sorry.

19 MR. EUGENE ZARASHAW: I mean, it still  
20 wouldn't help us retrieve the data or find it in the  
21 systems. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

25 [REDACTED]



HIGHLY CONFIDENTIAL

Page 71

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

[REDACTED]

[REDACTED]

SPECIAL MASTER GARRIE: Do you guys have a data diagram to explain the flows?

MR. EUGENE ZARASHAW: I don't know of one. I think this would be a data that would need to create or find data diagrams for every flow type or every product type because --

SPECIAL MASTER GARRIE: I don't mean every flow type so let me be specific.

I would just -- what about just basic one? Like where I'm struggling, right, is there's a ton of systems and I need to understand if they are -- if it is too complicated or as intertwined or as cumbersome as it sounds, there has to be some data diagrams that will explain or articulate how this user data flows.

I mean, because, like -- because you have DS share and -- I mean, you know, I could give you an example, like what websites a user visits? Like is that tracked?

MR. EUGENE ZARASHAW: I don't know offhand. I know that underlying data would arrive at Facebook servers just because of the browser always sending the HTTP refer header any time.

But I don't know if we actually report it

HIGHLY CONFIDENTIAL

Page 72

1 anymore.

2 I'm not sure if the HP refer header at this  
3 point is kept or dropped.

4 SPECIAL MASTER GARRIE: Mr. Elia, do you know?

5 MR. STEVEN ELIA: I also don't know. I did  
6 briefly work on the social plug-ins, things like the  
7 like button that you --

8 SPECIAL MASTER GARRIE: That's literally my  
9 next question.

10 MR. STEVEN ELIA: But during my time working  
11 on that, which was 2018 and 2019, I don't recall ever  
12 seeing that data get processed and used --

13 SPECIAL MASTER GARRIE: So my --

14 MR. STEVEN ELIA: -- or stored.

15 SPECIAL MASTER GARRIE: So here's my -- does  
16 Facebook receive data about user off-platform activity  
17 via share or their like plug-in?

18 MR. STEVEN ELIA: As Eugene said, it's  
19 certainly received. It's sent just as an ordinary part  
20 of browser behavior. But whether that results in any  
21 sort of logging or writing to any data source, I'm not  
22 aware of any, just from seeing that plug-in.

23 SPECIAL MASTER GARRIE: Do you have a data --  
24 there has to be some -- or takes custody of, may be a  
25 better way is we use the word take custody of the data.

HIGHLY CONFIDENTIAL

Page 73

1 Does Facebook take custody of the data when it  
2 comes inbound?

3 MR. EUGENE ZARASHAW: I don't really know what  
4 that term means.

5 SPECIAL MASTER GARRIE: So I -- maybe try it  
6 differently.

7 Is there a data diagram that would explain to  
8 me how Facebook handles share and like plug-in activity?

9 So I'm on another website. I click something  
10 or a like or a share.

11 Does Facebook take custody or receive or  
12 whatever that data, how -- like a data diagram  
13 explaining how the user of that data flows and through  
14 what systems?

15 MR. EUGENE ZARASHAW: I don't know of one.

16 Russ, do you happen to know -- is there one we  
17 already produced for this or one that exists?

18 MR. FALCONER: I'm also not aware of one off  
19 the top of my head. And I don't want to speak off the  
20 cuff about that, but that's something we can check. But  
21 nothing comes to mind right now.

22 SPECIAL MASTER GARRIE: Yeah. I mean, so I  
23 get the save part, but I know the C part is definitely  
24 inbound data off the header; right? So you saw it.

25 So the question is: Is that associated with

## HIGHLY CONFIDENTIAL

Page 74

1 the user? And then -- maybe that's the ad. I don't --  
2 I'm just trying to figure out -- 'cause you're getting  
3 that data, right, from a user's off-platform activity.  
4 It's coming inbound. You're capturing it.

5 What about -- I guess about what via partner,  
6 the API? So I'm a partner. I make Facebook an API and  
7 share the data back inbound about users.

8 MR. EUGENE ZARASHAW: So we do have an API for  
9 uploading graph activities via the SDK.

10 SPECIAL MASTER GARRIE: But what about via  
11 partner making? So a partner like an Amazon or  
12 somebody, or a Netflix makes an API available to you?

13 MR. EUGENE ZARASHAW: The only API like that  
14 I know of as -- that I'm recalling is the push  
15 notification API, the Android and iOS push  
16 notifications.

17 SPECIAL MASTER GARRIE: What about you,  
18 Mr. Elia?

19 MR. STEVEN ELIA: There is a similar kind of  
20 push functionality, for instance, in a platform product  
21 we have called WebHooks. The contents of those push  
22 updates are defined by Facebook. The developer would  
23 configure both, which types of events they want to be  
24 notified about, and a URL on their server to receive  
25 those.

HIGHLY CONFIDENTIAL

Page 75

1           SPECIAL MASTER GARRIE: So then the WebHooks  
2 is like a published subscribed standard; right?

3           So my question is -- I guess the way I think  
4 about it is does Facebook receive data about users'  
5 off-platform activity via a partner making an API  
6 available to Facebook?

7           So I'm Amazon. I make an API available that  
8 Facebook could call attached to the user because there's  
9 a whole lot of conversation about reciprocity.

10          MR. EUGENE ZARASHAW: I don't know of any such  
11 API.

12          MR. STEVEN ELIA: Same. I'm not aware of any.

13          SPECIAL MASTER GARRIE: So then what about  
14 partners bulk uploading data to Facebook? Not use --  
15 I'm not saying user specific. I'm just saying partners  
16 of bulk uploading the data into Facebook? Batch jobs.

17          MR. EUGENE ZARASHAW: So the only thing that  
18 would remotely be like that is some of the work on the  
19 ad side, but I'm not familiar with the details.

20          SPECIAL MASTER GARRIE: Okay. So then all the  
21 off-platform activity goes to pixel share or like  
22 basically, it sounds like.

23          Is there any other way that it's tracked, the  
24 off-platform activity?

25          MR. EUGENE ZARASHAW: It's possible for a --

HIGHLY CONFIDENTIAL

Page 76

1 an off -- a third party to upload a custom audience of  
2 identifiers for their own ad targeting, using our API or  
3 our website upload functionality to upload a list of  
4 either anonymized or non-anonymized identifiers to us,  
5 and then use that list to target their ads without  
6 knowing anything about the content of that list.

7 SPECIAL MASTER GARRIE: And you're hashing  
8 them like -- so do they know who you're serving ads to?

9 MR. EUGENE ZARASHAW: There are multiple ways  
10 to -- there are multiple types of identifiers that can  
11 be used. And I'm quite out of date in this. I was  
12 never particularly deep. But we had functionality to  
13 upload user IDs, and at one point we had functionality  
14 to target by e-mail addresses.

15 I'm not sure what other form -- types of  
16 targeting there were where a partner would just upload  
17 the set or -- there was also functionality for hash sets  
18 where we wouldn't even know the e-mail addresses that  
19 they're targeting directly, unless they were also  
20 interfaced with users, and then they could just use  
21 that custom audience to run their campaign against.

22 SPECIAL MASTER GARRIE: I know a couple other  
23 companies do that.

24 What about you, Mr. Elia?

25 MR. STEVEN ELIA: I'm not aware of any others.

HIGHLY CONFIDENTIAL

Page 77

1           SPECIAL MASTER GARRIE: How hard would it be  
2 for Facebook to take the UID and the pixel ad user or  
3 unit and associate -- and figure out what systems they  
4 actually have of the remaining systems or remaining  
5 43 systems?

6           MR. EUGENE ZARASHAW: It would take multiple  
7 teams on the ad side to track down exactly the -- where  
8 the data flows. I would be surprised if there's even a  
9 single person that can answer that narrow question  
10 conclusively.

11           I think part of the challenge we're running  
12 into is we tend to build pieces of infrastructure that  
13 are generic and to -- whenever possible. And then just  
14 leave them running for anybody at the company to use.

15           And sometimes piece of infrastructure end up  
16 using other pieces of infrastructure as underlying  
17 storage and ends up being as a layer on it. But they're  
18 not necessarily built with a clear set of only use  
19 cases.

20           So it's very difficult to let's say start with  
21 a system and find out all of the users of that system.

22           On the other -- it's somewhat simpler to start  
23 with specific use cases you have and find all the  
24 engineers involved in implementing that use case, and  
25 from that find out all the systems that might end up

## HIGHLY CONFIDENTIAL

Page 78

1 being used tracing through when you use a system to  
2 understand what other systems that system in turn uses.

3 SPECIAL MASTER GARRIE: I read your patents  
4 and there's definitely inferences about users, like  
5 their personality traits.

6 Where does that information go then?

7 MR. EUGENE ZARASHAW: For that one we would  
8 need someone from the ads team, probably from the ads  
9 targeting team.

10 SPECIAL MASTER GARRIE: Because, I mean,  
11 there's patents to predict all sorts of things about  
12 people that I assume are implemented somehow.

13 So if you're able to infer someone's rate of  
14 being married or not, which is in one of your patents,  
15 or anything -- you know, something of the like, that  
16 data is being inferred because I can't imagine users  
17 telling you "I think we're going to get divorced."

18 So there has to be somewhere -- at some level  
19 that data has to be aggregated out and inferred from the  
20 user activity, and then stored somewhere.

21 MR. EUGENE ZARASHAW: I just don't know  
22 offhand where we stored with -- it would take tracking  
23 down multiple people in the ads team to find out all of  
24 the systems that might store this.

25 SPECIAL MASTER GARRIE: I'm in the unfortunate



## HIGHLY CONFIDENTIAL

Page 79

1 situation where the Court has made it very clear that  
2 the information relating to a user -- when we say user  
3 name, plaintiffs is -- needs to be determined.

4 So then how does Facebook acquire data about  
5 user activity off-platform from third parties like  
6 Netflix or Amazon or others?

7 MR. EUGENE ZARASHAW: That's getting beyond my  
8 depth, if we do it all.

9 SPECIAL MASTER GARRIE: Mr. Elia?

10 MR. STEVEN ELIA: To the degree those other  
11 services use things like Facebook log-in, we would  
12 certainly know when a user has done so, logged in with  
13 their Facebook account, that is.

14 When a developer like Netflix, for instance,  
15 uses the user access token, they get as a result of that  
16 log in to make API requests. That is also something we  
17 would know about --

18 But I'm not aware of other instances outside  
19 of ordinary use of our APIs by those parties that might  
20 be what you have in mind.

21 SPECIAL MASTER GARRIE: Let's talk about just  
22 the APIs; right?

23 API. So we -- yesterday -- the data -- let's  
24 start with APIs.

25 I'm Netflix. There's an API that's made

## HIGHLY CONFIDENTIAL

Page 80

1 available. I mean, is that a -- we'll get into the  
2 private API conversation. But how is that mapped to the  
3 user? You said the token? Is that what you said?

4 MR. STEVEN ELIA: An access token is used to  
5 authenticate API requests.

6 Towards the beginning of our session today,  
7 Eugene talked about two different kinds of tokens. One  
8 used on things like mobile devices and web browsers that  
9 is specific to the user, and another type of token for  
10 server-to-server API requests.

11 And so the -- result of the user going through  
12 that log-in flow --

13 THE COURT REPORTER: Sorry. I'm not sure if I  
14 missed something. The Zoom froze. So what I heard at  
15 the end was "And so the result of the user going through  
16 that log-in flow."

17 MR. STEVEN ELIA: The developer gets a user  
18 token as a result of going through that log-in flow.

19 SPECIAL MASTER GARRIE: Is that user token the  
20 same thing as a user ID?

21 MR. STEVEN ELIA: It's not the same thing.  
22 But a user access token is a secure encryption of  
23 primarily the user ID and the app ID. So that when that  
24 token is provided with an API request, Facebook's APIs  
25 can understand which user and which app is making this

HIGHLY CONFIDENTIAL

Page 81

1 request.

2 SPECIAL MASTER GARRIE: So then how does --  
3 so then --

4 MR. EUGENE ZARASHAW: It's a bearer token that  
5 covers both authentication and authorization, and then  
6 access control is handled on the app server based on the  
7 author's -- based on the ID that has been verified  
8 inside the token. So the -- sorry.

9 SPECIAL MASTER GARRIE: No. Go ahead.

10 MR. EUGENE ZARASHAW: So the access control is  
11 based on the pair of user ID and app ID. The access  
12 token -- the bearer of the access token is authorized to  
13 act on behalf of the payer of the intersection of user  
14 and app. So the lesser of the privileges of the two.

15 SPECIAL MASTER GARRIE: So then could we use  
16 the access token with the user ID to find all the data  
17 for third-party apps at least that were accessing data  
18 across the face of these systems?

19 MR. EUGENE ZARASHAW: You could use it to  
20 query APIs for access to data that a third-party app  
21 would have about that user if you knew which exact data  
22 you were to query for.

23 So as an example, a complete synthetic  
24 example, let's assume an app had access to a photo. A  
25 user posted a photo and an app user payer had access to

HIGHLY CONFIDENTIAL

Page 82

1 the photo, you would need to know what photo to pull up  
2 or how to navigate through the Graph API to find the  
3 list of photos and then find the photo you're looking  
4 for.

5 SPECIAL MASTER GARRIE: So then let's return  
6 to a specific example for me. Let's talk about Netflix  
7 or Amazon. You can pick either one.

8 What data does Facebook get about a user's  
9 Netflix -- Amazon activity, and where does that data end  
10 up being stored?

11 MR. EUGENE ZARASHAW: I'm not confident on  
12 this one because I have not worked with Amazon, but my  
13 understanding is we get absolutely nothing about Amazon.

14 SPECIAL MASTER GARRIE: What about Netflix?

15 MR. EUGENE ZARASHAW: We would get data if  
16 they were to use the Javascript SDK about when the  
17 Javascript SDK was loaded. It is served over the main  
18 difference from Facebook dot com, if I remember  
19 correctly --

20 Steven, correct me if I'm wrong.

21 -- so we might not get all the cookies.

22 I don't know if they use the social plug-ins,  
23 such as like and share. And we would have to check what  
24 they actually send us via the Graph API.

25 SPECIAL MASTER GARRIE: Okay. So then,

HIGHLY CONFIDENTIAL

Page 83

1 Steven, anything -- Mr. Elia, anything you would like to  
2 add to that?

3 MR. STEVEN ELIA: No. My understanding is the  
4 same.

5 SPECIAL MASTER GARRIE: So then where does  
6 that data get stored? I mean, there has to be -- I  
7 remember the SDK. I just don't know on the back end  
8 where it goes.

9 MR. EUGENE ZARASHAW: [REDACTED]  
[REDACTED]

11 SPECIAL MASTER GARRIE: The SDK has  
12 function -- yeah. Keep going. I'm sorry, Mr. Zarashaw.

13 MR. EUGENE ZARASHAW: Logging about the raw  
14 HTTP request by default ends up in Hive.

15 If API calls are made, it would depend which  
16 API is being called and where the data would get stored  
17 so...

18 SPECIAL MASTER GARRIE: Do we have a data  
19 diagram for that? Like you develop -- someone must have  
20 a diagram that says this is where this data is stored.

21 MR. EUGENE ZARASHAW: We have a somewhat  
22 strange engineering culture compared to most where we  
23 don't generate a lot of artifacts during the engineering  
24 process. Effectively the code is its own design  
25 document often.

HIGHLY CONFIDENTIAL

Page 84

1           SPECIAL MASTER GARRIE: So you have to look at  
2 the code to figure out -- I mean, so --

3           MR. EUGENE ZARASHAW: For what it's worth,  
4 this is terrifying to me when I first joined as well.

5           SPECIAL MASTER GARRIE: You understand,  
6 though, like -- so you're telling me to figure this out,  
7 I have to look at the Netflix code to see how they're  
8 calling it, and then talk to each engineer in resource  
9 to figure out where it's actually storing that data that  
10 you're getting? There's no diagram?

11           Because you wrote an SDK so there has to be  
12 someone that knows what's supposed to go where at a  
13 minimum?

14           MR. EUGENE ZARASHAW: I don't think you'd have  
15 to talk to Netflix. It would be more -- someone would  
16 need to analyze what is Netflix actually uploading to  
17 us, which would be likely somewhere in Hive.

18           Then someone on the team that handles that  
19 specific API, because many teams at the company expose  
20 different APIs. The API stack itself is another piece  
21 of infrastructure that multiple teams contribute to.

22           So we would need to step on, find out what  
23 Net -- API's Netflix is calling that might result in  
24 data being exempt.

25           Step 2 --

HIGHLY CONFIDENTIAL

Page 85

1           SPECIAL MASTER GARRIE: So we have at least a  
2 high level understanding of that and the time period of  
3 that.

4           So where does it go?

5           MR. EUGENE ZARASHAW: So Step 2 would need to  
6 find the engineering teams responsible for the specific  
7 APIs and ask them to track down where are all of the  
8 places the data ends up?

9           And it would depend on exactly what is being  
10 up loaded and what the API is for, whether it ends up --  
11 like if you're looking at -- we can certainly give you a  
12 probabilistic answer. But to get a certain answer would  
13 take some digging.

14          SPECIAL MASTER GARRIE: All right. But --  
15 okay. Probabilistic.

16          What about you, Mr. Elia? I mean, you have  
17 engineers on your team and I get the fact that you set  
18 them loose on the world. But like there's an SDK for an  
19 API. So there's APIs. They have SDK. There's a  
20 Javascript SDK. I program it. It comes through to your  
21 system. Someone somewhere must know where that data is  
22 going to -- it goes to the Hive.

23          Then you -- what happens there? It gets  
24 extracted out of the Hive and stored into other systems  
25 or how does this work?

HIGHLY CONFIDENTIAL

Page 86

1 MR. STEVEN ELIA: So to use that example, that  
2 SDK could call any number of APIs, and our code base is  
3 open where any engineer in the company would have the  
4 ability to modify that SDK, build a new API for it to  
5 call.

6 And so as Eugene was saying, I think that  
7 would require identifying all the APIs, someone like  
8 Netflix was calling, finding the teams that are  
9 responsible for those APIs.

10 And last I checked last year there were  
11 thousands of APIs and over 500 teams that had built  
12 these.

13 And so the individuals from the teams  
14 responsible for those APIs would then be the ones most  
15 knowledgeable about --

16 SPECIAL MASTER GARRIE: I get it.

17 But we could seduce statistically  
18 significantly sample and come up with a relatively  
19 probabilistic model that would answer probably for  
20 95 percent of the data that flows through the company.  
21 I mean, you're not talking that -- I mean, we don't need  
22 to look at a thousand. I could look at 20 of them and  
23 aggregate that through to all of them with a 95 percent  
24 confidence level at some level.

25 But what I don't understand is -- so basically



## HIGHLY CONFIDENTIAL

Page 87

1     you don't make it -- to engineer up front is required to  
2     make data flow diagrams or data diagrams or anything  
3     about what data is going where?

4             MR. EUGENE ZARASHAW: That is the exception of  
5     the rule.

6             SPECIAL MASTER GARRIE: True, Mr. Elia?

7             MR. STEVEN ELIA: Yes. That's correct.

8             SPECIAL MASTER GARRIE: So we have 43 systems  
9     where engineers of a --

10            MR. EUGENE ZARASHAW: For what it's worth, in  
11     the design of a lot of these systems, there are diagrams  
12     and design artifacts for how the underlying piece of  
13     infrastructure is designed.

14            However, it is rare for there to exist  
15     artifacts and diagrams on how those systems are then  
16     used and what data actually flows through them. They're  
17     generally agnostic to the data.

18            MR. STEVEN ELIA: I would add onto that that,  
19     yes, often these systems may be built with a particular  
20     use case in mind. If other teams find that they have  
21     their own use case that would benefit from using that  
22     system, they're typically free to make use of it.

23            And so even the teams who maintain and are  
24     responsible for these systems may not be aware of all  
25     of the specific use cases that are leveraging.

HIGHLY CONFIDENTIAL

Page 88

1           SPECIAL MASTER GARRIE: Well, with that in  
2 mind, then the DIY file is certainly incomplete.  
3 Because nobody actually knows what the users are  
4 actually doing.

5           I mean -- I mean, that would be the logical  
6 conclusion.

7           All right. Let me ask the next question then.  
8 Sorry for thinking out loud there.

9           So then is there aggregated and inferred data  
10 that is extracted from the Hive relevant to the users,  
11 being the named plaintiffs here, on other named -- you  
12 know, that's my first question, and then I have a  
13 follow-up question to that.

14          MR. EUGENE ZARASHAW: What do you mean by  
15 relevant in this case?

16          SPECIAL MASTER GARRIE: I mean anything. They  
17 have brown cows, red cows, green cows, purple cows. I  
18 mean, anything. Like your Hive literally stores  
19 everything people do in the entire whole wide existence.

20          So I have no idea what some -- based on what  
21 you told me, some Facebook engineer can write a query  
22 and be, like, do these people have brown cows and  
23 associated with their user ID because it's related to  
24 the cow liking button.

25          I mean, is there -- so, I mean, any -- is

## HIGHLY CONFIDENTIAL

Page 89

1     there any aggregated or inferred data that's extracted  
2     from the Hive relevant to a user?

3             MR. EUGENE ZARASHAW: I don't know. Hive is  
4     another similar system where it's relatively free form,  
5     how it can be used.

6             So it's back to the starting with the final  
7     artifact and trying to trace out how the artifact was  
8     created is much easier than starting with the original  
9     data and trying to trace all the artifacts that may have  
10    been created from it.

11            SPECIAL MASTER GARRIE: So you can't  
12    aggregate -- but can a developer aggregate and extract  
13    data out of the Hive that is aggregated or inferred?

14            MR. EUGENE ZARASHAW: They can write queries  
15    to -- assuming they pass the privacy checks, which does  
16    require getting certain levels of permission and have a  
17    good reason to be doing it, they can create data based  
18    on other data in Hive during the data and so on and  
19    intermediate tables.

20            But there's no way to find out, for example,  
21    what are all of the intermediate tables a user appears  
22    in.

23            MR. FALCONER: Could I ask -- I'm sorry. I  
24    just want to make sure I'm following.

25            When you say a developer,

HIGHLY CONFIDENTIAL

Page 90

1 Special Master Garrie, do you mean like an internal  
2 Facebook developer or like a --

3 SPECIAL MASTER GARRIE: Yeah. Only a Facebook  
4 engineer. I understand only Facebook engineers have  
5 direct access to the Hive.

6 Is that not true?

7 MR. FALCONER: Right. Okay.

8 MR. EUGENE ZARASHAW: That's correct.

9 SPECIAL MASTER GARRIE: Okay.

10 So I've sat through a lot of depositions so I just  
11 assume certain things, but fair point.

12 So let me give you a more practical example.

13 Can an engineer extract linguistic data from  
14 the Hive that has communications between different  
15 users, retrieve those data from the user profile; right?  
16 So -- because the messaging is part of the DIY file.

17 So I extract linguistic data from the DIY  
18 file, I then retrieve characteristics from the user's  
19 profile themselves; right? So we have linguistic and  
20 the user profile, and I'm extracting data from both of  
21 them.

22 And then I determine, based on the user's  
23 linguistic patterns and their profile data, I then store  
24 certain personality or characteristics traits of it;  
25 right?

## HIGHLY CONFIDENTIAL

Page 91

1 I know you guys filed for patents on this.

2 My question is where does that data go?

3 Because then you store the content for the user based  
4 on those characteristics that you derive, at least  
5 according to the patents.

6 MR. EUGENE ZARASHAW: Since we're on the  
7 record, I don't believe that case is possible because to  
8 my recollection, the message contents don't end up in  
9 Hive, but in their own separate storage.

10 SPECIAL MASTER GARRIE: Yeah. Sorry. I  
11 forgot. You're right.

12 But you can look -- so you look at the message  
13 and the other -- so you look at the message and the user  
14 profile, not in Hive. Sorry. So remove it from Hive.

15 You look at them outside of Hive?

16 MR. EUGENE ZARASHAW: So it would depend where  
17 that team wanted to store it. There are multiple pieces  
18 of infrastructure they could use. They could also build  
19 custom infrastructure if they wanted to store something  
20 like that.

21 It would entirely depend on the specific team  
22 in question, what they chose to use.

23 SPECIAL MASTER GARRIE: So how I do find out  
24 which systems actually have this user data that can be  
25 inferred?

HIGHLY CONFIDENTIAL

Page 92

1 MR. EUGENE ZARASHAW: I don't know.

2 SPECIAL MASTER GARRIE: Seems very socialized.

3 MR. EUGENE ZARASHAW: It's a rather difficult  
4 conundrum.

5 SPECIAL MASTER GARRIE: I mean, I appreciate  
6 the equality of access here. It's just -- it seems  
7 very -- I'm trying to figure out a way that's more  
8 efficient and effective to meet what the Court has  
9 ordered for the 43 other systems. But it doesn't sound  
10 like -- let's keep going.

11 So what other databases housed -- do you  
12 have -- is there any list anywhere the databases that  
13 housed inferred or aggregated data? Is there -- of the  
14 systems we have, can you tell me which ones are likely  
15 to house it?

16 MR. EUGENE ZARASHAW: I don't know. Hive is  
17 the only one is my go-to guess. Probabilistically.

18 SPECIAL MASTER GARRIE: And Hive is one of the  
19 42 that remain.

20 Like, Mr. Elia, do you have any guesses?

21 MR. STEVEN ELIA: I don't have an indication  
22 about some of the others. I probably used half a dozen  
23 of these myself.

24 SPECIAL MASTER GARRIE: Which ones do you use?  
25 Let's focus on those.

## HIGHLY CONFIDENTIAL

Page 93

1 MR. STEVEN ELIA: I'll go through this list.  
2 EverStore, Memcache, TAO, UDB and XDB,  
3 MySQL, although I mostly think of that in terms of UDB  
4 and XDB.

5 And those are the ones that I have personal  
6 familiarity with.

7 SPECIAL MASTER GARRIE: Then could TAO store  
8 that information?

9 MR. STEVEN ELIA: Could you repeat --

10 SPECIAL MASTER GARRIE: The inferences?

11 MR. STEVEN ELIA: Could you repeat the full  
12 context of your question?

13 SPECIAL MASTER GARRIE: Yeah.

14 So I look at the chat message -- so you look  
15 at my chat messages with a friend, and then you look at  
16 my profiles and you decide that I am -- you make some  
17 inferences about my personality traits or  
18 characteristics or -- and then store that in a database.

19 MR. STEVEN ELIA: TAO might be an odd choice  
20 to use for that. What you could do if you were to store  
21 information like that in TAO is create a new kind of  
22 object with fields for these different forms of  
23 inferences you wanted to keep track of and associate  
24 that object, for instance, with the user those are  
25 inferences about.

HIGHLY CONFIDENTIAL

Page 94

1 SPECIAL MASTER GARRIE: But it sounds like  
2 MySQL could be just as well a database if they wanted,  
3 they could do it in?

4 MR. STEVEN ELIA: It could. That, as Eugene  
5 also mentioned, is typically not used for user data,  
6 just because it wouldn't support it at that scale.

7 SPECIAL MASTER GARRIE: Okay. So that -- I  
8 would assume EverStore would not be one of those?

9 MR. STEVEN ELIA: EverStore is used to store  
10 things like the photo or video binary files. It's more  
11 of like a blob kind of storage.

12 I also, like Eugene, don't know if it even has  
13 user IDs. So given something like an identifier of a  
14 photo that's in EverStore, I don't know if you could  
15 even find the user who uploaded that photo.

16 SPECIAL MASTER GARRIE: But it is linked  
17 somehow to the user ID, otherwise you just have a bunch  
18 of photos sitting --

19 MR. STEVEN ELIA: [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

24 SPECIAL MASTER GARRIE: So then TAO is sort of  
25 like the database on top of the database, so to speak.



HIGHLY CONFIDENTIAL

Page 95

1 I have pictures --

2 MR. EUGENE ZARASHAW: It's often used that  
3 way.

4 SPECIAL MASTER GARRIE: So then you look at  
5 TAO.

6 And then if you -- then what about --  
7 what's -- you mentioned Mem?

8 MR. STEVEN ELIA: Memcache, which is a caching  
9 layer, so during a request there might be some  
10 information that's repeatedly accessed.

11 SPECIAL MASTER GARRIE: That's the ADD  
12 experience. I got it.

13 User satisfaction, eminent -- or most -- user  
14 experience.

15 MR. STEVEN ELIA: As a -- let me give a  
16 concrete example.

17 So something like your own name might appear  
18 multiple times in the course of loading a particular  
19 Facebook page like your profile. Rather than query TAO  
20 in some underlying database every time that's needed to  
21 render, that would be expensive. It's cheaper to use a  
22 caching layer like Memcache.

23 SPECIAL MASTER GARRIE: Yeah. That's why they  
24 have ram -- yeah. Same idea, just a software extracted  
25 ram. I get it. It's for that user experience. People

## HIGHLY CONFIDENTIAL

Page 96

1       aren't patient. I get it. And it's expensive. And  
2       your data centers and all that.

3               Then you said UDB, and what was the other one  
4       you mentioned?

5               MR. STEVEN ELIA:   XDB.

6               SPECIAL MASTER GARRIE:   XDB.

7               Would they be used to store that?

8               MR. STEVEN ELIA:   UDB and XDB, to the degree  
9       there's differences between them, I'm not familiar with  
10      them. I would use those almost synonymously. These are  
11      often direct MySQL instances.

12              My own familiarity and use of them has been  
13      with the capability system that I described yesterday in  
14      my testimony for whitelisting apps for a certain kind of  
15      functionality.

16              SPECIAL MASTER GARRIE:   Yeah. I thought  
17      that's where all that sort of data rests or existed;  
18      right?

19              MR. STEVEN ELIA:   There's an XDB instance that  
20      stores that information. But any team might have their  
21      own XDB instance for data specific to their product or  
22      system.

23              SPECIAL MASTER GARRIE:   There's no way to  
24      query the system to find out if there's user data  
25      inferences in it, for UDB or XDB, it sounds like?

HIGHLY CONFIDENTIAL

Page 97

1 MR. EUGENE ZARASHAW: Not that I know of.

2 MR. STEVEN ELIA: Yeah. I don't think at a  
3 technical level it would even be possible to query  
4 across all UDB or XDB instances.

5 SPECIAL MASTER GARRIE: You have to write a  
6 script and then type it out and dump it, and then query  
7 that script. And then go back and look at that dataset.

8 MR. EUGENE ZARASHAW: It's often more painful  
9 than a script because you're often dealing with hundreds  
10 of thousands of instances. So even one failure can  
11 cause issues.

12 It takes some nontrivial engineer to work  
13 through reliably, even touch every single database  
14 shard.

15 SPECIAL MASTER GARRIE: Yeah. I mean, again,  
16 we could sample it and skip every and just do a sample.

17 Is whitelist, you mean in and out -- an allow  
18 list; right? That's what you -- when you say -- the  
19 whitelist is a capability -- yesterday there was a lot  
20 of -- in your prior testimony, there was discussion  
21 about capabilities and whitelists and that tool set.

22 That's where that set is; right? In an  
23 instance in that database?

24 MR. STEVEN ELIA: The list of capabilities  
25 that have been granted to a particular application and

HIGHLY CONFIDENTIAL

Page 98

1 the history of that is in an XDB, as well as -- as of  
2 last year, I believe, that information is also  
3 duplicated in TAO.

4 SPECIAL MASTER GARRIE: Looks like TAO data is  
5 pretty -- it's fairly -- what systems in here on this  
6 list of 55 -- or I'll call it 43 -- are TAOs an  
7 overlayer for most user data? Do you have any idea?

8 MR. STEVEN ELIA: I don't know which of these  
9 systems TAO would duplicate info of.

10 [REDACTED]  
[REDACTED]  
[REDACTED]  
[REDACTED]

14 SPECIAL MASTER GARRIE: A primary -- okay.

15 So what other data -- is it possible then to  
16 tell me what other data, other than anonymized data, is  
17 in the database or in these systems here?

18 If I were to say I don't care about anonymized  
19 data, I just want to know about non-anonymized data in  
20 the systems, is there any way to figure out?

21 MR. EUGENE ZARASHAW: I don't know of one.  
22 We're back to square one on having to figure out what is  
23 actually stored in each database and often task the  
24 engineering team that owns that specific data what it  
25 actually is.

HIGHLY CONFIDENTIAL

Page 99

1 SPECIAL MASTER GARRIE: Okay.

2 MR. EUGENE ZARASHAW: For what it's worth,  
3 there are some systems we could rule out as definitely  
4 being duplicate or transient data that is -- like, for  
5 example, Memcache is a good example of only storing  
6 things that are, okay, oozing in short order because  
7 it's just caching data.

8 Or Unicorn is another one where it's just an  
9 indexing system where the primary data must live  
10 somewhere else.

11 Laser is another system where by definition,  
12 any data it serves is just fronting an existing Hive  
13 table.

14 So there are certain systems that -- it's very  
15 difficult to find out what has user data but it might be  
16 possibly to rule out some systems as having data always  
17 derived from some other system.

18 SPECIAL MASTER GARRIE: Yeah. That's what it  
19 sound -- it sounds like some of these -- or they have  
20 data that's contained in another system; right?

21 So TAO --

22 MR. EUGENE ZARASHAW: There's a lot of  
23 duplication because the use cases for the different  
24 systems vary more often by the read rather than write  
25 case, since that tends to be the higher volume.

## HIGHLY CONFIDENTIAL

Page 100

1           So there are often systems that exist simply  
2       because the system they cover can't handle the read  
3       volumes.

4           SPECIAL MASTER GARRIE: Do either one of you  
5       have any idea what OIL is?

6           MR. EUGENE ZARASHAW: No.

7           MR. STEVEN ELIA: I don't.

8           SPECIAL MASTER GARRIE: It would be good to  
9       know what the acronyms stand for at a minimum for the  
10      systems.

11          Some I know, like -- what about Poseidon?

12          MR. EUGENE ZARASHAW: Unfortunately not.

13          MR. STEVEN ELIA: I'm not familiar with it.

14          SPECIAL MASTER GARRIE: What's TW Storage?

15          MR. EUGENE ZARASHAW: Don't know.

16          MR. STEVEN ELIA: I also don't know what it  
17      is. I don't know with confidence what "TW" stands for,  
18      but that may be Tupperware, which is a common way we  
19      would abbreviate a system named that.

20          SPECIAL MASTER GARRIE: I think we're at  
21      another breaking point. If you want, we can take a  
22      break, and then I have -- yes. I do have more  
23      questions. And we'll take a five-minute break and then  
24      I'll let -- well, if Facebook has any questions, and  
25      then plaintiffs have any questions, we'll ask them right

HIGHLY CONFIDENTIAL

Page 101

1 after the five-minute break.

2 Hopefully -- oh, Mr. Espinoza is here so I  
3 won't drop anyone this time. And I apologize, Mr. Elia.  
4 That is likely me.

5 So we will take a five-minute break. We'll go  
6 off the record and then we'll resume.

7 We'll go off the record.

8 (Break taken in proceedings.)

9 SPECIAL MASTER GARRIE: All right to resume  
10 for the Facebook side?

11 MS. WEAVER: We do have some questions, if  
12 we're allowed.

13 JAMS CONCIERGE CARLOS ESPINOZA: Mr. Falconer,  
14 just trying logging off and logging back in again. I'll  
15 see if that does the trick to reset your mic.

16 SPECIAL MASTER GARRIE: I think he's -- we're  
17 still off the record so...

18 (Pause in proceedings.)

19 SPECIAL MASTER GARRIE: In the interest of  
20 keeping everything on track, we're going move this right  
21 along here. We're good to go back on the record?

22 THE COURT REPORTER: Ready.

23 SPECIAL MASTER GARRIE: All right. So  
24 plaintiffs, what questions would you like to ask?

25 MS. WEAVER: I think we still are interested

HIGHLY CONFIDENTIAL

Page 102

1 in learning what the replacement ID is, its relationship  
2 to the Facebook ID. And if Facebook is going to satisfy  
3 an individual user's data deletion request, how does it  
4 accomplish that? And does it trace the data through  
5 some form of IDs or a script or a query?

6 SPECIAL MASTER GARRIE: That's my GDPR line of  
7 questioning I have. So the replacement ID, I'm willing  
8 to entertain.

9 Is there -- does Facebook use something called  
10 a replacement ID?

11 MR. EUGENE ZARASHAW: I don't know what that  
12 is.

13 Steven?

14 MR. STEVEN ELIA: I've heard of an RID, but I  
15 don't know if that's the same thing.

16 SPECIAL MASTER GARRIE: That answers that. It  
17 would be good if Facebook could check -- well, actually  
18 before that happens, plaintiffs, can you clarify where  
19 did you learn about a replacement ID?

20 MS. WEAVER: It's been referenced in some of  
21 the documents produced.

22 I can actually -- if Chris Springer wants to  
23 identify it, I don't know if he knows or not.

24 But we also know that there are other  
25 alternative -- and I guess we think the replacement ID



HIGHLY CONFIDENTIAL

Page 103

1 is the RID. So if you can explain the context of what  
2 the RID is?

3 We're just trying to trace the identifiers  
4 that can be used to identify content that can be linked  
5 under the definition of personal information in the  
6 State of California that can be reasonably associated  
7 with these named plaintiffs.

8 What are the identifiers that we use to do  
9 that?

10 MR. DEREK LOESER: To provide slightly more  
11 context, there are documents produced that describe the  
12 deletion process of de-anonymizing data, and that's  
13 related to the same query.

14 SPECIAL MASTER GARRIE: I have a full line of  
15 questioning around the GDPR, but -- or CCPA.

16 I guess we can ask those questions first and  
17 then we'll progress forward.

18 For GDPR, or the European Innovation Statute  
19 or the California Consumer Privacy Act, they have very  
20 stringent requirements around users' right to request  
21 information about them to be deleted.

22 How does Facebook actually define a user and  
23 delete the data?

24 Before you answer, what I'm specifically  
25 interested in is do you have diagrams or any

HIGHLY CONFIDENTIAL

Page 104

1 documentation that's set out how that works and  
2 operates?

3 MR. EUGENE ZARASHAW: I would have to defer to  
4 the deletion and privacy team on this one. It is  
5 definitely not my area of expertise.

6 MR. FALCONER: Yeah. That -- I would suggest  
7 here that this is a pretty special team in -- at the  
8 company that works on this and is knowledgeable about  
9 this. And obviously we're more than happy to answer  
10 these questions, but I want to get the right people to  
11 give you the best answers.

12 SPECIAL MASTER GARRIE: That's fine.

13 MR. FALCONER: And so if we could do that, I  
14 think that might be more fruitful.

15 SPECIAL MASTER GARRIE: Not a problem.

16 I want to talk to the engineering resources  
17 that worked with Mr. Pope on that effort. I believe he  
18 was a technical project manager in that effort. And so  
19 it would be good to just -- the engineering resources  
20 for how they did it because they surveyed all of these  
21 people to make this list.

22 So we'll table that for a follow on, maybe.

23 Okay. RID, can you clarify what's the RID and  
24 the chance that is, indeed, the replacement ID?

25 MR. STEVEN ELIA: You're directing that to me?

HIGHLY CONFIDENTIAL

Page 105

1 SPECIAL MASTER GARRIE: Yeah, Mr. Elia.

2 No lawyer is my question directed at at all.

3 MR. STEVEN ELIA: So I'll give you my limited  
4 understanding of it. I think this would also be a topic  
5 that could be useful to discuss with an engineering team  
6 responsible for it.

7 But I've seen RIDs described as a way to  
8 anonymize data.

9 A common scenario would be to [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

23 SPECIAL MASTER GARRIE: That's an interesting  
24 way to -- does that work across all of the system --  
25 what systems is that RID operationally applied to?

HIGHLY CONFIDENTIAL

Page 106

1 MR. STEVEN ELIA: I'm not sure which systems  
2 outside of TAO might store data off of an RID.

3 SPECIAL MASTER GARRIE: All right. So next  
4 sort of bucket of things we're going to talk about is  
5 third-party integration APIs.

6 I know Facebook required -- at least  
7 documentation I've reviewed, allows third-party  
8 applications to access private APIs that are not  
9 publicly available.

10 Is that your understanding as well, I guess,  
11 Steven or Eugene?

12 MR. STEVEN ELIA: That's right.

13 MR. EUGENE ZARASHAW: Yes.

14 SPECIAL MASTER GARRIE: So now I want to ask  
15 very specific technical questions because I've been  
16 curious about this.

17 What do the following permissions provide:  
18 Friends and anyone?

19 MR. STEVEN ELIA: The friends --

20 SPECIAL MASTER GARRIE: Friends underscore of  
21 underscore anyone.

22 MR. STEVEN ELIA: Yes. Friends of anyone is a  
23 capability that allows an application who would have it  
24 to view friendships without the users in question  
25 needing to have installed or logged into that

HIGHLY CONFIDENTIAL

Page 107

1 application.

2 Ordinarily the public behavior of this is that  
3 for a friendship to be visible, both users need to have  
4 logged into that app and both need to have authorized  
5 access to their friends list.

6 And so this capability allows doing that even  
7 if, say, one of those users had not logged into that  
8 application.

9 SPECIAL MASTER GARRIE: Okay. And then what  
10 about the XMPP underscore log-in underscore grant?

11 MR. STEVEN ELIA: I'm not familiar with that  
12 one off the top of my head.

13 MR. EUGENE ZARASHAW: I am.

14 X -- the XMPP protocol is an old chat  
15 protocol. And Facebook used to operate a set of chat  
16 servers that I believe could integrate with XMPP client  
17 a very long time ago.

18 So this permission has to do with the ability  
19 to log into the chat servers back when it was chat,  
20 before it was even called Messenger, from what I  
21 remember.

22 I do not know if anybody still has this.

23 SPECIAL MASTER GARRIE: What about -- but why  
24 would you use it? What would you use it for?

25 MR. EUGENE ZARASHAW: I'm now getting a little

HIGHLY CONFIDENTIAL

Page 108

1 hazy on my memory, but if I recall, it's to make sure  
2 only certain clients could log into the chat servers,  
3 the Facebook chat servers.

4 SPECIAL MASTER GARRIE: Okay. What about  
5 TLS without dialogue?

6 MR. EUGENE ZARASHAW: Oh, yes. I do remember  
7 that one.

8 And, Steven, correct me if I'm wrong on this  
9 one.

10 I believe this was for apps which build their  
11 own grant permissions dialogue UI, usually embedded apps  
12 where they could not display a web browser with our UI  
13 on them. So instead we would have a contract where they  
14 would specify exactly how they were going render the UI  
15 with all the disclaimers we needed.

16 And in exchange, we would let them grant app  
17 permissions after log-in without displaying our  
18 dialogue.

19 SPECIAL MASTER GARRIE: So there's no user  
20 consent obtained from the Facebook side. It was relying  
21 on the app provider?

22 MR. EUGENE ZARASHAW: Exactly.

23 And the type of consent, the exact pixels for  
24 the consent, would be spelled out in the contract. So  
25 that all of the consent would be in the third-party

HIGHLY CONFIDENTIAL

Page 109

1 controlled UI, but they would be contractually obligated  
2 to still show everything we needed them to show.

3 SPECIAL MASTER GARRIE: I saw that on all  
4 your -- a lot of the telco mobile...

5 MR. EUGENE ZARASHAW: Exactly.

6 SPECIAL MASTER GARRIE: Because it comes  
7 pre-installed I guess.

8 So then standard friends info, friends of  
9 anyone, I didn't get how that was different.

10 MR. STEVEN ELIA: If I recall correctly --  
11 it's been a few years since I've looked at these  
12 friends-related capabilities -- it's very similar to  
13 the friends of anyone use case.

14 One of them, I'm not sure which, is restricted  
15 to the type of access tokens and requests made from a  
16 developer's servers as opposed to things like the mobile  
17 device or web browser.

18 SPECIAL MASTER GARRIE: So then -- and then  
19 what about FQL friend request permission with auto  
20 accept link?

21 MR. STEVEN ELIA: Could you repeat the name?

22 SPECIAL MASTER GARRIE: Yeah. FQL underscore  
23 friend underscore request underscore permission  
24 underscore with underscore auto accept underscore link.

25 MR. STEVEN ELIA: I believe this was a

HIGHLY CONFIDENTIAL

Page 110

1 capability which allowed the app to see the friend  
2 requests a user has received and information about those  
3 friend requests, like who it's from, as well as a link  
4 which they could render to the user, and if the user  
5 clicks that link, it would accept that friend request.

6 SPECIAL MASTER GARRIE: Okay. I get it.

7 So what about nearby underscore users?

8 MR. STEVEN ELIA: I don't recall off the top  
9 of my head what that capability did.

10 MR. EUGENE ZARASHAW: I don't know this one  
11 either.

12 SPECIAL MASTER GARRIE: What about Titan  
13 underscore API?

14 MR. STEVEN ELIA: That was a capability that  
15 was generally used to gain access to messaging-related  
16 APIs. So an API that would allow sending a message or  
17 reading messages that a user had received, for example.

18 SPECIAL MASTER GARRIE: Now, these APIs could  
19 write to which database -- which of those systems we've  
20 been -- those 55 systems?

21 MR. STEVEN ELIA: The most common that they  
22 would write to would be TAO or Hive. Not being familiar  
23 with most of these 55, I can't say which others. But at  
24 a technical level, any of these that have the ability to  
25 be written to from production traffic are technically



HIGHLY CONFIDENTIAL

Page 111

1 possible.

2 SPECIAL MASTER GARRIE: So then if we call  
3 these -- let's say if we called TOS without dialogue  
4 capability, what -- that would write to TAO? Like how  
5 does this capability -- like what I'm trying to  
6 understand is capabilities -- who defines where the  
7 data writes to on a capability?

8 MR. STEVEN ELIA: A capability is used to  
9 gate a section of code. It could be an entire API. It  
10 could be a particular field on an API. It could be a  
11 particular function called by an API.

12 And so the capability itself doesn't directly  
13 result in any kind of data to be written, but determines  
14 whether a given piece of code should execute based on  
15 whether the app making the request has that capability.

16 SPECIAL MASTER GARRIE: So then where does the  
17 data for the TOS get written?

18 MR. EUGENE ZARASHAW: There are two sets of  
19 data with that specific capability. One is the --  
20 effectively the metadata for an app itself.

21 SPECIAL MASTER GARRIE: Right.

22 MR. EUGENE ZARASHAW: And the fact that  
23 Facebook has granted this specific application ID, this  
24 capability, which would be something on the contracts,  
25 review process and so on, and that would likely live in

HIGHLY CONFIDENTIAL

Page 112

1 an XPP.

2 Steven, correct me if I'm wrong --

3 MR. STEVEN ELIA: That's right. That's right.

4 MR. EUGENE ZARASHAW: And might also be cached  
5 in TAO.

6 Then the other thing would be where does the  
7 data get written that is the result of that -- of an  
8 API call being used where the behavior of the API call  
9 is different because of a capability?

10 So for that specific capability --

11 And, again, Steven, correct me if I'm wrong.

12 -- the relevant API call would be the grant  
13 permissions dialogue or the -- or auth dot log in where  
14 when the user logs into the app, because the capability  
15 is present, some additional side effects could happen.  
16 And those side effects, because it's about the  
17 relationship between the user and the app, would likely  
18 be stored in TAO.

19 SPECIAL MASTER GARRIE: Okay. So basically  
20 these are all the capabilities that are needed for  
21 building apps to do what you guys do today?

22 MR. EUGENE ZARASHAW: Most of the capabilities  
23 around, yes, building a -- allowing a third party to  
24 build all or part of a replacement Facebook client  
25 themselves.

## HIGHLY CONFIDENTIAL

Page 113

1 SPECIAL MASTER GARRIE: Does Facebook have a  
2 record of all the private APIs?

3 MR. EUGENE ZARASHAW: Go ahead, Steven.

4 MR. STEVEN ELIA: There's a record of APIs  
5 which are in their entirety gated by a capability,  
6 rendering it a private API.

7 To the degree an API is publicly available but  
8 might have slightly different behavior for an app with a  
9 particular capability, that would be harder to produce.  
10 And I can't think of an explicit record of that.

11 SPECIAL MASTER GARRIE: Is the capability what  
12 makes the API itself private then?

13 MR. STEVEN ELIA: If the API is entirely gated  
14 behind that capability, yes, the capability is what  
15 makes it a private API.

16 SPECIAL MASTER GARRIE: So it essentially can  
17 only get the capability via contract basically, it  
18 sounds like, or in an agreement or something?

19 MR. STEVEN ELIA: These days we have a whole  
20 review process. I talked about this in yesterday's  
21 testimony called Partner Grant Review for granting  
22 capability steps.

23 SPECIAL MASTER GARRIE: But 2014, pre-2014?

24 MR. STEVEN ELIA: Pre-2014, I'm less familiar  
25 with. That would generally have been managed by the

HIGHLY CONFIDENTIAL

Page 114

1 partnership's organization.

2 SPECIAL MASTER GARRIE: But there was no  
3 capability API private wall guard, as you're talking  
4 about?

5 MR. STEVEN ELIA: There were capabilities  
6 pre-2014. There was not a tool that made it easy to  
7 view current capability grants or the capabilities a  
8 particular app has before 2014.

9 SPECIAL MASTER GARRIE: So then what about  
10 figuring out -- is there any documentation for -- or  
11 some documentation, or any -- let's say any  
12 documentation for what user data the private APIs could  
13 access? I mean, any. That's caps, "ANY".

14 MR. STEVEN ELIA: This is an instance where a  
15 phrase Eugene used before I think is relevant. We would  
16 often think of the code itself as the documentation.

17 So for a given private API, it would be very  
18 explicit what fields could be returned or requested as  
19 part of that API, and that would be in the code for that  
20 API.

21 SPECIAL MASTER GARRIE: The comment in the  
22 code at the top of the code file list?

23 MR. STEVEN ELIA: It's not in a comment, but  
24 the structure of an API's code has a very explicit  
25 delineation of the fields that can make up a response to

HIGHLY CONFIDENTIAL

Page 115

1     that API.

2                 MR. EUGENE ZARASHAW: It went so far as we  
3     were able to automatically generate documentation by  
4     using some descriptors in the code.

5                 SPECIAL MASTER GARRIE: Because if it would --  
6     so then you could automatically generate that  
7     information for the private APIs if they all follow the  
8     same -- if the program --

9                 MR. EUGENE ZARASHAW: If it becomes -- well,  
10    that would cover the APIs. It wouldn't cover the exact  
11    behaviors of what happens when they are called, and the  
12    differences in behavior if a certain capability is  
13    present.

14                SPECIAL MASTER GARRIE: Yeah. But you would  
15    know if it opens an SQL connection to a database and  
16    writes to the database at some level.

17                At some point you have to open or stream a  
18    connection to write information or receive information;  
19    right?

20                MR. EUGENE ZARASHAW: Right.

21                It would just be a rather difficult task to  
22    automatically find out what exactly a deep code stack  
23    does.

24                SPECIAL MASTER GARRIE: That I get.

25                But the way the APIs are written, it sounds

HIGHLY CONFIDENTIAL

Page 116

1 like that you could function, at least identify the data  
2 elements.

3 MR. EUGENE ZARASHAW: It's clear which ones  
4 are reads and which ones are writes.

5 SPECIAL MASTER GARRIE: Yeah.

6 And do they -- and it defines what those data  
7 elements are in the code somehow?

8 MR. EUGENE ZARASHAW: It defines what is the  
9 write somebody is trying to make. As an example --  
10 let's use an example, the stream publishing API.

11 It defines as the write that a post may have a  
12 body, a link, a photo, and so on. And that's the end of  
13 the definition.

14 What happens to that data once it passes into  
15 the API code depends on exactly how that specific API is  
16 implemented, and therefore where it's stored.

17 So, for example, if you're uploading a photo,  
18 it may -- the photo may go into [REDACTED]

[REDACTED]

[REDACTED]

21 It is not clear from the API description or  
22 implementation up front where that storage actually is.  
23 You have to dig through many layers of code to  
24 understand which storage a specific write would end up  
25 in.

HIGHLY CONFIDENTIAL

Page 117

1           SPECIAL MASTER GARRIE: So basically there's  
2 no documentation?

3           MR. EUGENE ZARASHAW: There's documentation on  
4 the APIs but not what they actually do underneath.

5           There's documentation also of what is the  
6 product implication of the API. So there's  
7 documentation that if I call an API to post to feed, a  
8 story will appear in the news feed of the person who  
9 owned the access token containing what was posted.

10           But there's not documentation on exactly how  
11 that is implemented underneath.

12           SPECIAL MASTER GARRIE: So then I guess my  
13 next question is: Does the DIY tool include all the  
14 data available that's accessed via the private APIs?

15           I assume no. Definitely not.

16           But anyways, that's my question.

17           MR. EUGENE ZARASHAW: I would defer to Russ on  
18 what exactly does include, since I don't know much about  
19 the DIY tool.

20           MR. FALCONER: Yeah. And let us get back with  
21 a definitive answer on that.

22           But our understanding, which I think we've  
23 submitted before, is that the DYI has -- whatever data  
24 might be called up through calling any API end point  
25 would be included, you know, in the most comprehensive,

HIGHLY CONFIDENTIAL

Page 118

1 most human readable form in the DYI.

2 That's my recollection. But, again, let us  
3 get you a formal confirmation on that.

4 SPECIAL MASTER GARRIE: Yeah. But I'm talking  
5 about -- let's be clear. I'm talking about the private  
6 APIs where Netflix or somebody else is using a private  
7 API to generate data, I would assume, since there's no  
8 documentation, that the DIY file -- they wouldn't even  
9 know possibly to include it in the DYI file.

10 Is that right, Mr. Zarashaw?

11 MR. EUGENE ZARASHAW: I don't know that that's  
12 a reasonable assumption just because there's very little  
13 documentation anywhere, and yet we have a DYI file that  
14 contains quite a lot.

15 SPECIAL MASTER GARRIE: Yeah. It's a mystery  
16 to me as to how the DYI file is actually built and then  
17 where does all the -- because all of the inference data  
18 that you can derive about a user's behavior isn't stored  
19 in the DYI file. It's stored and distributed in other  
20 datasets because you don't -- Facebook doesn't provide  
21 inferences that it derives about user behavior in the  
22 DIY file.

23 MR. EUGENE ZARASHAW: To my knowledge, the  
24 reason it works is many, many years of an effort. It  
25 was not a simple project.



HIGHLY CONFIDENTIAL

Page 119

1           SPECIAL MASTER GARRIE: Well, that's my  
2 question --

3           MR. FALCONER: And, again, we're happy to  
4 verify this formally. But there are -- our  
5 understanding is there are derived inferences of a  
6 DYI file pretty voluminously.

7           SPECIAL MASTER GARRIE: That would be good --  
8 I mean, yeah, so that would be good to understand.

9           And then which ones aren't. Because it  
10 doesn't sound like there's documentation anywhere.

11           So what I'm trying to figure out is since we  
12 don't know what actually could be written to any of the  
13 systems relating to the user data, how do we know that  
14 the data and the user DIY file is possibly even close to  
15 complete if there's no recorded documentation?

16           But we'll talk about that offline. I don't  
17 want to waste the engineer's time.

18           But that's just my question; right? I mean,  
19 because the Court is very clear about it being, you  
20 know, data collected from the user's on-platform  
21 activity, data obtained from a third party regarding a  
22 user's off-platform activities, and data inferred from  
23 the user's on-and-off platform activity.

24           And it's clearly stated that you collect in  
25 your terms -- at least in your interrogatories, that you

HIGHLY CONFIDENTIAL

Page 120

1 collect the Bluetooth, the wifi signal, and all that  
2 through the mobile and associate it with the user. But  
3 that doesn't -- as far as I can tell, my review of the  
4 DIY file didn't actually have that information, just as  
5 an example. Or, you know, any of the other data in  
6 there.

7 But you can -- because what I'm trying to  
8 understand is there are 40 -- or 55 total systems, but  
9 we can't seem to figure out what's actually stored from  
10 a user data perspective in the systems because the  
11 developer could store whatever they -- data wherever  
12 they wanted as they saw fit.

13 MR. EUGENE ZARASHAW: Mr. Garrie, the solution  
14 to this is unfortunately exactly the work that was done  
15 to create the DYI file itself.

16 And the thing I struggle with here is in order  
17 to find gaps in what may not be in DYI file, you would  
18 by definition need to do even more work than was done to  
19 generate the DYI files in the first place.

20 And if it helps, we can find out approximately  
21 the size of the team and the number of months it took  
22 that team to do the work to sift through the systems and  
23 generate DYI file as a scoping exercise for what it  
24 would take to find even more, if there's more to find.

25 SPECIAL MASTER GARRIE: Well, it sounds like

HIGHLY CONFIDENTIAL

Page 121

1 TAO is a good place to start. It sounds like to me TAO  
2 and maybe, like, one or two other systems are a good  
3 place to rest your laurels to go from, based on my  
4 review of the patents and the conversations so far.

5 But I guess my next question is are user  
6 message threads -- they are -- are the user message  
7 threads in the DIY tool accessed via the social graph?

8 MR. EUGENE ZARASHAW: If the social graph  
9 contains the backing store for the messages product,  
10 which is one point called Titan, and I don't recall what  
11 the system is called, then yes.

12 SPECIAL MASTER GARRIE: And then let's be  
13 clear. Like, I mean, to deal with the EU, there has to  
14 be some process about the data and the deletion of data  
15 and user data at some level to be compliant with GDPR,  
16 whether it's inferential data or not inferential data.  
17 So we'll follow up with Russ on that separately.

18 So then --

19 MR. EUGENE ZARASHAW: My understanding of the  
20 deletion process for, to use a good analogy, you can  
21 throw away a key even if you don't know where the lock  
22 is.

23 SPECIAL MASTER GARRIE: Yeah. All right.  
24 That's one way to do it.

25 MR. EUGENE ZARASHAW: It's by far the easier

## HIGHLY CONFIDENTIAL

Page 122

1 way to do it. There's a very similar problem of do you  
2 sift through a massive box --

3 SPECIAL MASTER GARRIE: I get it. You got the  
4 point across. I get it. I get it.

5 I'm telling you now, as long as the database  
6 stays and isn't stolen. But yeah.

7 So then do private APIs provide data other  
8 than the data in the social graph?

9 MR. EUGENE ZARASHAW: Not to my knowledge, in  
10 large part because the -- most of the other systems  
11 could not handle the load.

12 SPECIAL MASTER GARRIE: Steven -- or Mr. Elia?

13 MR. STEVEN ELIA: No. I think almost by  
14 definition user data that is part of the Graph API,  
15 including any private API, is to be part of the social  
16 graph.

17 SPECIAL MASTER GARRIE: All of them or just  
18 some of them?

19 MR. STEVEN ELIA: Some of the -- any API,  
20 including the private API, that would, for example,  
21 return user data, I would classify that as part of the  
22 social graph.

23 SPECIAL MASTER GARRIE: What about inferred  
24 user data? So not user input data but user-inferred  
25 data. So I'm drawing a distinction; right?

HIGHLY CONFIDENTIAL

Page 123

1           There's user data you input and then there's  
2           user data you infer. Two different things.

3           MR. STEVEN ELIA: Sure.

4           I'm not aware of Graph APIs that return  
5           inferred data.

6           SPECIAL MASTER GARRIE: So where does all that  
7           go?

8           MR. STEVEN ELIA: Other than the very  
9           high-level understanding of things like our ads  
10          targeting system leveraging that kind of data, I'm not  
11          sure what uses inferred data like that.

12          SPECIAL MASTER GARRIE: Because that would be  
13          on-platform activity data about a user's behavior that  
14          is being inferred and stored and captured by Facebook.

15          MR. STEVEN ELIA: Besides use cases like our  
16          own ads targeting, I'm not aware any of that includes  
17          anything that would share that with a third-party  
18          developer.

19          SPECIAL MASTER GARRIE: Well, not -- it's not  
20          a third-party developer. It's if it -- it's if it's  
21          on -- so -- so that would make -- the Court defines it  
22          as data collected from a user's on-platform activity,  
23          not data shared with a third party. It's -- let's be  
24          clear.

25          The delimiter isn't whether it's shared. The

HIGHLY CONFIDENTIAL

Page 124

1 delimiter is data collected from a user's on-platform  
2 activity.

3 MR. STEVEN ELIA: I apologize. I was  
4 mentioning that because of the connection to private  
5 APIs, which I'm considering to be --

6 SPECIAL MASTER GARRIE: Yeah.

7 MR. STEVEN ELIA: -- APIs called by third  
8 parties.

9 SPECIAL MASTER GARRIE: Yeah. Fair enough.  
10 But I just want to make the distinction here;  
11 right? This is where it becomes more challenging,  
12 right, because it's data collected from a user's  
13 on-platform activity so you have on-platform messages  
14 and you have on-platform user profile and you draw  
15 inferences about that user's behavior and store that in  
16 the database. We're not sure which database systems.  
17 Maybe TAO, maybe Timeline, like a handful of systems  
18 for -- that's group 1.

19 And then you have data obtained from third  
20 parties regarding a user's off-platform activities.

21 That is your APIs; right? That is your data  
22 dumps and your private APIs. It doesn't sound like you  
23 do data dumps but APIs at a minimum.

24 And then there is data inferred from a user's  
25 on-or-off platform activity. Again, not defining what

HIGHLY CONFIDENTIAL

Page 125

1     inferred is, but it says, you know, on-or-off platform  
2     activity.

3             So if you get data returned from Netflix and  
4     you make -- you make inferences based on their movie  
5     watching habit and their chat messages and you write  
6     that out to a third party -- one of your internal  
7     database systems. So the API covers really the second  
8     bucket, not buckets 1 or 3 that the Court's defined.

9             But we'll keep going.

10            All right. Partner integration. And then  
11     we'll take a break, and then I think Counsel Falconer  
12     will be closer to the finish line.

13            So partner integration, what access does --  
14     'cause you define in the -- in multiple different  
15     places, you define strategic partners, such as Amazon,  
16     Etsy, Netflix, et cetera.

17            What data did they get for that relevant time  
18     period?

19            MR. STEVEN ELIA: Each of those would be its  
20     own unique case and depends on what is the specific user  
21     experience they were building. And the agreements, the  
22     contracts we had with them would outline what kind of  
23     data would be made available as part of that and what  
24     they would be allowed to do with that.

25            I'm not that familiar with all of those

HIGHLY CONFIDENTIAL

Page 126

1 specific agreements.

2 SPECIAL MASTER GARRIE: What's a strategic  
3 primary you are familiar with? We can work backwards  
4 from that.

5 MR. STEVEN ELIA: Bing is an example of one  
6 which I've implemented some of the APIs they would use.  
7 Bing the search engine.

8 SPECIAL MASTER GARRIE: Yeah. So that's good  
9 because that's something I know something about.

10 All right. So Bing and Facebook. You can  
11 search friends of friends on Bing; right?

12 How did that work?

13 MR. STEVEN ELIA: I'm not aware of the ability  
14 to search that on Bing.

15 What I am aware of is an API we had for them  
16 that shared public content on Facebook. So this were  
17 things like post photos, videos that had public privacy,  
18 anyone on Facebook could see it, and those kind of posts  
19 could show up if they were relevant to your search  
20 query, for instance.

21 SPECIAL MASTER GARRIE: And did Facebook get  
22 any data back from Bing?

23 MR. STEVEN ELIA: I'm only aware of a unit  
24 directional API here of Facebook pushing data to Bing.

25 SPECIAL MASTER GARRIE: Are there any



HIGHLY CONFIDENTIAL

Page 127

1       bidirectional partner APIs?

2               MR. STEVEN ELIA: I can't think of any off the  
3       top of my head.

4               SPECIAL MASTER GARRIE: Where would one look  
5       to find that out? The engineer, not -- not just --

6               MR. STEVEN ELIA: Sure.

7               So a rather laborious exercise I could carry  
8       out as an engineer, for instance, would be to start with  
9       a list of all the capabilities, including ones that have  
10      ever existed.

11              For each one, look to see whether it gated an  
12      API. It would be much more difficult to see did it  
13      modify behavior of an otherwise public API?

14              And in some cases I might have to go through  
15      the entirety of the code history for that API or frankly  
16      other files in the repository that may have referenced  
17      that capability to see what that did.

18              I don't think even in that case that that  
19      would yield any insight into an API the partner provided  
20      that Facebook would have called.

21              SPECIAL MASTER GARRIE: What about looking at  
22      the firewall logs for the inbound data?

23              MR. STEVEN ELIA: I'm not familiar with that  
24      part of our infrastructure and what's possible there.

25              MR. EUGENE ZARASHAW: The one challenge I

HIGHLY CONFIDENTIAL

Page 128

1 recall from looking at the inbound logs is that we did  
2 not store post-request bodies, and most write APIs were  
3 post requests.

4 SPECIAL MASTER GARRIE: That is very true.

5 Okay. Did -- and it sounds like -- and I did  
6 look that the A -- private -- well, the strategic  
7 partners had varying network access, additional  
8 computing and storage.

9 Is that accurate?

10 MR. STEVEN ELIA: [REDACTED]  
[REDACTED]

12 SPECIAL MASTER GARRIE: Yeah. That's it.

13 MR. STEVEN ELIA: [REDACTED]  
[REDACTED]  
[REDACTED]  
[REDACTED]

17 SPECIAL MASTER GARRIE: Did Facebook allow  
18 data connections between the Facebook and its strategic  
19 partners?

20 MR. EUGENE ZARASHAW: What do you mean by  
21 data connections in this case?

22 SPECIAL MASTER GARRIE: Did they allow  
23 partners to connect to Facebook's data stores?

24 MR. EUGENE ZARASHAW: No. All access was  
25 through the API.

HIGHLY CONFIDENTIAL

Page 129

1 SPECIAL MASTER GARRIE: So scheduled API  
2 calls?

3 MR. EUGENE ZARASHAW: Would WebHooks count?

4 SPECIAL MASTER GARRIE: Yeah, I guess.  
5 What about FTP, SFTP servers?

6 MR. EUGENE ZARASHAW: Not to my knowledge.  
7 That's terrifying.

8 MR. STEVEN ELIA: I'm also not aware of any  
9 instances like that.

10 SPECIAL MASTER GARRIE: Is there a difference  
11 between custom APIs and private APIs?

12 MR. STEVEN ELIA: I wouldn't use the phrase  
13 "custom API" myself. That doesn't mean anything to me.  
14 Maybe given a specific context sentence it was used in,  
15 I'd be better able to answer that.

16 SPECIAL MASTER GARRIE: That didn't mean  
17 anything to me either so -- I see it a lot, so I was  
18 just wondering if Facebook created a term called  
19 "custom APIs."

20 So I assume then there are query-able Hive use  
21 cases for data that exists now.

22 Is that a -- is that accurate?

23 MR. EUGENE ZARASHAW: Do you mean for API  
24 usage data?

25 SPECIAL MASTER GARRIE: Or for debugging.

HIGHLY CONFIDENTIAL

Page 130

1 MR. EUGENE ZARASHAW: There's a fair bit of  
2 data in [REDACTED]  
3 [REDACTED], and so on.

4 SPECIAL MASTER GARRIE: Right.

5 So there are other tools for querying.

6 Is there a query-able Hive use cases today for  
7 user data?

8 MR. EUGENE ZARASHAW: I'm not sure what you  
9 mean by use cases.

10 There are Hive tables that can be queried  
11 using our standard query UI. And if you know what table  
12 to look in and it has the right fields, you can query  
13 for user data.

14 SPECIAL MASTER GARRIE: Is there a list of --  
15 or --

16 So the answer is yes, you have the use --  
17 there are use cases for where you're querying high data  
18 tables?

19 MR. EUGENE ZARASHAW: Yes.

20 SPECIAL MASTER GARRIE: Could you -- how --

21 Sorry. I'm just thinking.

22 So then where was data acquired by a partner  
23 retrieved from?

24 MR. EUGENE ZARASHAW: It would depend on the  
25 API but most APIs are backed by TAO.

HIGHLY CONFIDENTIAL

Page 131

1 SPECIAL MASTER GARRIE: What about you,  
2 Mr. Elia, do you agree?

3 MR. STEVEN ELIA: I agree with that statement.

4 MR. EUGENE ZARASHAW: There are some system  
5 specific APIs that for a given set -- or certain -- some  
6 specific APIs for product would -- may be backed by a  
7 different back end.

8 So as an example, messaging has its own  
9 separate back end, and any of those Titan APIs you  
10 mentioned earlier would be partly backed by [REDACTED]  
[REDACTED]

12 Another example would be reading the news feed  
13 for a user, reading the stories they have posted [REDACTED]  
[REDACTED] [REDACTED]  
[REDACTED]

16 SPECIAL MASTER GARRIE: Where is the message  
17 back end?

18 MR. EUGENE ZARASHAW: I don't recall the new  
19 name of it, but I remember it being -- looking up some  
20 of the 55 systems, the message back end was on that  
21 list.

22 SPECIAL MASTER GARRIE: Okay. It would be  
23 good to identify which one that is.

24 And then the user data and the message back  
25 end I assume stores way more data than what sits in the

HIGHLY CONFIDENTIAL

Page 132

1     DIY file.

2                 MR. EUGENE ZARASHAW: I'm not sure what's in  
3     the DIY file, so --

4                 SPECIAL MASTER GARRIE: DIY, do it yourself.  
5     Sorry.

6                 MR. EUGENE ZARASHAW: No, no. I meant I'm not  
7     sure what's in that file. I'm not sure what it actually  
8     contains versus what's in the back end.

9                 SPECIAL MASTER GARRIE: All right.

10                MR. EUGENE ZARASHAW: To answer your earlier  
11     question, the name of the current messenger back hand is  
12     Callisto. So it is No. 4 on that 55 item list.

13                SPECIAL MASTER GARRIE: Okay. DYI files.

14                So then can a system exist that queries the  
15     Hive for data to create inferences and store it  
16     elsewhere?

17                MR. EUGENE ZARASHAW: Yes.

18                SPECIAL MASTER GARRIE: Do you know what  
19     systems that those inferences are being written up to?

20                MR. EUGENE ZARASHAW: I don't know a system  
21     like that could exist. I don't know if it does. But if  
22     one were to exist, I would likely ask the ad targeting  
23     team if they have something like that.

24                SPECIAL MASTER GARRIE: Okay. You said end  
25     targeting team.

HIGHLY CONFIDENTIAL

Page 133

1 MR. EUGENE ZARASHAW: Ads targeting. Yes.

2 SPECIAL MASTER GARRIE: Okay. Did you say

3 "ends" or "end"?

4 MR. FALCONER: Ads, a-d-s.

5 MR. EUGENE ZARASHAW: Ads.

6 MR. STEVEN ELIA: Advertisements.

7 SPECIAL MASTER GARRIE: Oh, advertising  
8 targeting team. Got it. Ads. That makes more sense.

9 Sorry. Go ahead.

10 MR. EUGENE ZARASHAW: I was going to say a  
11 common way of storing derived data from Hive is also in  
12 Hive.

13 SPECIAL MASTER GARRIE: But you could look for  
14 code that talks to the Hive?

15 MR. EUGENE ZARASHAW: [REDACTED]

[REDACTED]  
17 SPECIAL MASTER GARRIE: Yeah. But your  
18 production code being the consumer face -- you use --  
19 you say production, that's the consumer facing. That's  
20 not your research, your development, or your  
21 engineering.

22 MR. EUGENE ZARASHAW: Generally for any kind  
23 of randomized access, for example, if we want to serve  
24 ads, the data would have to be pulled out of Hive into  
25 another system for immediate data access. Hive has very

HIGHLY CONFIDENTIAL

Page 134

1 high latency.

2 SPECIAL MASTER GARRIE: Oh.

3 MR. EUGENE ZARASHAW: In the minutes.

4 So one example system on that list is No. 9,  
5 Laser, which does not store any data of its own but it  
6 is a cache in front of other systems that are slow and  
7 can't handle the load, such as can be TAO, can be Hive,  
8 and so on, to accelerate them. So it's an in-memory  
9 cache in front of a slower system.

10 SPECIAL MASTER GARRIE: That would be very  
11 helpful, I guess, Counsel Falconer, if you can identify  
12 those systems that are caching systems versus systems  
13 that actually have beyond just user experience.

14 MR. FALCONER: Mm-hm. Yeah. Understood. We  
15 can do that.

16 SPECIAL MASTER GARRIE: Okay. So then did  
17 integration involve Facebook -- so then is there any  
18 centralized area where -- you would receive data from a  
19 partner. We've established there's no way to  
20 determine -- for a particular user; right? I log in  
21 with whatever information, you would get that data back  
22 from a partner.

23 Do you know where that data would be stored?

24 I'll give you an example.

25 I go to Netflix, I use my Facebook ID. You



HIGHLY CONFIDENTIAL

Page 135

1 get back a set of movies I like. I don't believe the  
2 set of movies I like that Netflix shares with Facebook  
3 sits in my DIY file about me.

4 MR. EUGENE ZARASHAW: So the one way I could  
5 think that working is there's an API for posting that a  
6 user has watched something. It's the graph -- it's the  
7 graph actions API.

8 So at one point I believe Netflix may have  
9 been using it. I don't recall for sure. Steven might  
10 know. But --

11 A site -- if some site out there provided  
12 videos that people could watch and that site integrated  
13 with a Facebook API and made post requests to this API  
14 every time a user who had tossed the app, watched a  
15 movie, we would likely store that information in two  
16 places in duplicate: [REDACTED]

[REDACTED]  
18 [REDACTED]  
[REDACTED]  
[REDACTED]  
[REDACTED]  
[REDACTED]

22 SPECIAL MASTER GARRIE: I got it.

23 No reflection on this hearing.

24 Mr. Elia, is there anything you'd like to add  
25 to the examples?

HIGHLY CONFIDENTIAL

Page 136

1 MR. STEVEN ELIA: Nope.

2 SPECIAL MASTER GARRIE: And then -- so then do  
3 we -- is there any way to figure out from this list of  
4 systems which ones stored data that was accessed and/or  
5 received by partners, especially in cases where access  
6 was not through the Graph API? So --

7 MR. EUGENE ZARASHAW: I don't know of any ways  
8 for partners to get data outside of the API.

9 SPECIAL MASTER GARRIE: The Graph API or any  
10 API?

11 MR. EUGENE ZARASHAW: I don't -- I think it's  
12 all Graph API now.

13 Steven, correct me if I'm wrong.

14 We used to have a rest API that we folded into  
15 it but it was still the same exact API with slightly  
16 different parameters.

17 MR. STEVEN ELIA: There may be a small set of  
18 APIs that aren't based on the Graph API even today, but  
19 those would be well understood -- identified. We would  
20 have a list of those.

21 It's plausible that there could exist  
22 something not on that list.

23 But as a general matter of policy, again,  
24 there's a whole review process to go through whenever  
25 you're building this kind of API that would result in

HIGHLY CONFIDENTIAL

Page 137

1 such an API winding up in this list.

2 SPECIAL MASTER GARRIE: And pre-2014, that  
3 didn't exist?

4 MR. STEVEN ELIA: This type of review process  
5 I mention was even later that was started in 2018, 2019,  
6 something like that.

7 SPECIAL MASTER GARRIE: All right. And the  
8 Graph API that's now consolidated was not consolidated  
9 to a single API pre-20- -- what --

10 When did the graph API become the omnibus API  
11 for the partner access and receive?

12 MR. EUGENE ZARASHAW: Go ahead, Steven.

13 MR. STEVEN ELIA: For the general populous --  
14 '15/2016 was when we removed access to some of those  
15 other APIs, like FQL. But there, I believe, were  
16 capabilities allowing access to those for a few more  
17 years.

18 MR. EUGENE ZARASHAW: To add to that, the  
19 previous set of APIs also known as the rest APIs still  
20 followed the same semantics of needing an app ID,  
21 needing permissions, needing access tokens.

22 It was just different semantics and how to  
23 call them.

24 MR. STEVEN ELIA: I would also add to that  
25 that I'm not aware of any data, and I'm highly confident

## HIGHLY CONFIDENTIAL

Page 138

1 that the only data available through that legacy rest  
2 API FQL was also available over the Graph API. Those  
3 would be like duplicate APIs.

4 SPECIAL MASTER GARRIE: Who is the engineer  
5 that built the DYI tool?

6 MR. EUGENE ZARASHAW: I would have to defer to  
7 Russ on this one. It would be an entire team.

8 SPECIAL MASTER GARRIE: Okay. But on every  
9 team there's an engineer or two engineers that build --  
10 that do -- that are the pillars of the team.

11 MR. FALCONER: I -- go ahead. Sorry.

12 MR. EUGENE ZARASHAW: I was going to say, it's  
13 a more than two-pillar engineer undertaking.

14 MR. FALCONER: Yes. And that -- I don't -- we  
15 don't have names or statistics, you know, at our  
16 fingertips on that. We would be happy to get them.

17 But we have -- my understanding is the same as  
18 Mr. Zarashaw's about the scope of the project and how  
19 many people were involved.

20 SPECIAL MASTER GARRIE: Yes. So 30 engineers  
21 to one pillar? So a 60 engineer team?

22 MR. EUGENE ZARASHAW: I would be impressed if  
23 a 60 engineer team could have completed DYI tool in  
24 six months.

25 SPECIAL MASTER GARRIE: Not six months.

HIGHLY CONFIDENTIAL

Page 139

1 Figure 18 months. But yeah. And then they iterated on  
2 it for several years.

3 But anyways -- so specialized consumer  
4 experiences, are they different -- they're defined  
5 differently. Are they the same, though, specialized  
6 consumer experiences?

7 MR. EUGENE ZARASHAW: I'm not sure what that  
8 means.

9 SPECIAL MASTER GARRIE: So as I understood it,  
10 specialized consumer experiences, according to the  
11 Interrogatory 4, page 374:

12 "Facebook developed private APIs to  
13 enable select partners to offer custom  
14 seamless experiences for users seeking to  
15 more closely integrate their Facebook  
16 experiences with other things they'd like  
17 to do, such as listening to music,  
18 watching movies, or pursuing and sharing  
19 interests and hobbies. These custom  
20 experiences are built by companies such as  
21 DropBox, Spotify, and Netflix."

22 MR. EUGENE ZARASHAW: So it sounds like that  
23 refers to the whole set of private APIs where we allowed  
24 third parties to do customized integration into some  
25 aspects of how their experience worked with ours.

HIGHLY CONFIDENTIAL

Page 140

1           As an example, I recall we did special  
2 integration for DropBox that you mentioned where when  
3 you were sharing a DropBox folder with someone, you  
4 could also send them a message on Facebook about it,  
5 where for almost anyone else, there was no ability to  
6 send a message via an API but we enabled it via contract  
7 for this specific use case.

8           SPECIAL MASTER GARRIE: And theoretically that  
9 would be related to Calypso today?

10          MR. EUGENE ZARASHAW: Yes.

11          SPECIAL MASTER GARRIE: So it's either in the  
12 contract or the code. All right.

13          All right. So then I had some I guess DYI  
14 questions but we'll ask those later.

15          Measurement Partners, that's another term  
16 that's been defined as having different access, and  
17 it's -- where is it? -- do you know what  
18 Measurement Partners are?

19          MR. EUGENE ZARASHAW: I loosely know what they  
20 are but this is not my area of expertise.

21          SPECIAL MASTER GARRIE: Mr. Elia?

22          MR. STEVEN ELIA: Similar. I wouldn't have  
23 worked on any of their APIs.

24          SPECIAL MASTER GARRIE: They had their own set  
25 of APIs; right?

HIGHLY CONFIDENTIAL

Page 141

1 MR. EUGENE ZARASHAW: I don't know.

2 SPECIAL MASTER GARRIE: Mr. Elia?

3 MR. STEVEN ELIA: I don't specifically know.

4 If that was represented as an example of those  
5 customized integrations, then I would imagine that  
6 would be true, yes.

7 SPECIAL MASTER GARRIE: Do you know who was on  
8 the Measurement Partner engineering team, or who to talk  
9 to or --

10 MR. STEVEN ELIA: I don't know.

11 MR. EUGENE ZARASHAW: We would have to follow  
12 up to find who's actually on the team.

13 SPECIAL MASTER GARRIE: And the reason I ask  
14 is because there's a specific distinction and I'm trying  
15 to trace this back to these systems that exist because  
16 it says in your exhibit -- you know, in your  
17 interrogatory or -- anyways, the point is that there's  
18 references to individually identifiable user content  
19 that Facebook provides to the Measurement Partners  
20 that's stored.

21 And what's not clear to me is if that same  
22 information exists in the DYI file, and if it doesn't  
23 exist in the DYI file, where does it exist? Because you  
24 make it a point to state that you have this data and  
25 you're providing it.

HIGHLY CONFIDENTIAL

Page 142

1           And if it comes -- and my point is -- and this  
2   is worth noting, Counsel, for Facebook generally, if  
3   it's duplicative, that's great. But someone has to say  
4   it is duplicative because nobody is saying it's  
5   duplicative.

6           MR. FALCONER: Yeah. And, again, this is  
7   starting to veer more into the guts of the ads business,  
8   the advertising business. And so I think if this is a  
9   subject matter that the special master wants to ask more  
10   questions on, I think it would be better served  
11   getting -- putting you in touch with different people  
12   who were more familiar with that part of the business.

13           SPECIAL MASTER GARRIE: I'm not interested in  
14   the ad part of it. I'm interested in the fact that --  
15   and I can point you to what triggered it. It's an  
16   exhibit now exhibit in the depo as 128. It's page 495,  
17   and 128, page 374. And there's very explicit references  
18   to user, like individual identifiable -- which we've  
19   decided is not a term -- individually identifiable user  
20   content that Facebook provides to Measurement Partners.

21           What isn't stated is where it's stored, what  
22   is collected, and if it's duplicative of the DYI files.

23           And that's the same for the -- these SCEs or  
24   these specialized consumer experiences. And it's not  
25   clear to me -- and it's going to be important to



## HIGHLY CONFIDENTIAL

Page 143

1 understand what user data is collected in -- for these  
2 experiences that you've identified -- Facebook has  
3 identified, is it duplicative or is it in other one of  
4 these systems? Because there's no identification about  
5 where this data exists but it's clearly data that is  
6 created or collected from the user's on-platform  
7 activity.

8 MR. FALCONER: Sure. Yeah. Understood on all  
9 that.

10 And, again -- yeah. So let us find the right  
11 person -- I understand what you're interested in, and  
12 let us find the right person who can talk with you about  
13 that.

14 SPECIAL MASTER GARRIE: Then for user mobile  
15 device data, it sounds like we'll need a different set  
16 of engineers to provide that, unless it's completely  
17 duplicative of what's in the DYI file.

18 But, again, that's not established, and I  
19 can't figure out from the systems you're identifying  
20 what is or is not within the DYI file, specifically  
21 around the mobile user data. Because they go to  
22 Facebook's mobile app using Facebook engagement.

23 So, again -- but I think it sounds like,  
24 Mr. Zarashaw and Mr. Elia, you're not in a position to  
25 answer the questions about the mobile data?

HIGHLY CONFIDENTIAL

Page 144

1 MR. EUGENE ZARASHAW: I'm not.

2 MR. STEVEN ELIA: Are you referring to data,  
3 for example, collected from our mobile apps?

4 SPECIAL MASTER GARRIE: Yeah. That is part of  
5 the user.

6 MR. STEVEN ELIA: That's also again something  
7 I'm not familiar with. I wouldn't believe that, though,  
8 to be stored any differently from data collected  
9 generally from the web or elsewhere. I wouldn't think  
10 there's a separate, for example, data store of only  
11 mobile data, for instance.

12 SPECIAL MASTER GARRIE: Yeah. But there's  
13 different mobile -- there's different data that can be  
14 collected off mobile devices, and if it's part of the  
15 social graph and available, about the user activity.

16 MR. EUGENE ZARASHAW: I think the challenge  
17 I'm having is I don't actually know what's in the DYI  
18 file and I don't know what mobile data is collected. So  
19 it's difficult to do more than guess on what is stored  
20 for which.

21 SPECIAL MASTER GARRIE: Yeah. I mean, I could  
22 tell you in the terms of services that were provided by  
23 Facebook, it indicates that the Bluetooth ID, the  
24 network ID, a bunch of ID is collected.

25 So, I mean, again, Exhibit 128, page 487 kind

HIGHLY CONFIDENTIAL

Page 145

1 of delineates a bunch of that, just to point Facebook  
2 sort of in the direction of what I'm asking about or why  
3 I'm asking about it.

4 Because I did look at the files and I did not  
5 see these data elements in there, so I don't know if  
6 it's possible they didn't access it through their mobile  
7 devices or mobile -- or maybe they -- you know, but this  
8 may end up being is I'm just curious to understand that  
9 piece of it.

10 And then the last set of questions, we can  
11 either power through them or we can take a break.

12 What would Mr. Elia or Mr. Zarashaw prefer?

13 MR. EUGENE ZARASHAW: I would love a break.

14 MR. STEVEN ELIA: Same.

15 SPECIAL MASTER GARRIE: All right. That  
16 sounds good. I mean, the DYI file, as far as I could  
17 tell, has device ID, mobile servers provider, and  
18 country code. But not the other elements that are  
19 identified.

20 So we will take a -- you guys want to take a  
21 five-minute break?

22 Is that a yes, Mr. Elia.

23 MR. STEVEN ELIA: Yes.

24 SPECIAL MASTER GARRIE: We'll take a  
25 five-minute break. I'm supposed to put everybody in the

HIGHLY CONFIDENTIAL

Page 146

1 breakout room so that's probably -- we'll go off the  
2 record.

3 (Break taken in proceedings.)

4 SPECIAL MASTER GARRIE: All right. We're  
5 going to go back on the record.

6 I'll take, like we have, if there are  
7 questions for plaintiffs, and then I will jump into  
8 the --

9 Back on the record?

10 Plaintiffs, are there any questions you would  
11 like to ask?

12 MS. WEAVER: Yes.

13 Can you query the Hive with the Facebook user  
14 ID or some other identifier, and have they done that for  
15 the named plaintiffs, given the latency issue that  
16 Mr. Zarashaw mentioned?

17 SPECIAL MASTER GARRIE: I mean, the engineers  
18 don't really answer the second part of the question, but  
19 the first part of the question is can -- can you query  
20 the Hive for user data? I think the user that we  
21 already said is yes.

22 Is that correct, Mr. Zarashaw?

23 MR. EUGENE ZARASHAW: [REDACTED]

[REDACTED]

[REDACTED]

HIGHLY CONFIDENTIAL

Page 147

1

2

3

4

5

6

7

8

9

10 SPECIAL MASTER GARRIE: It would be possible,  
11 though, to identify the largest tables within the Hive  
12 and how those data schemas were; right? The top 10  
13 tables or whatever it may be.

14 MR. EUGENE ZARASHAW: Yes.

15 MS. WEAVER: May I ask a follow-up?

16 SPECIAL MASTER GARRIE: You may ask a question  
17 and then I will --

18 MS. WEAVER: Fine.

19 Is there a schema or fields for Hive or  
20 descriptors so that we can identify the top 10 or some  
21 way to identify what we could query using user ID?

22 SPECIAL MASTER GARRIE: I will reask the  
23 question.

24 If we could identify the top 10 tables, can  
25 you extract the fields for those top 10 tables

HIGHLY CONFIDENTIAL

Page 148

1     because -- whatever -- X number of tables. That's  
2     possible; right?

3             MR. EUGENE ZARASHAW: It is possible. But the  
4     only exception is going to be that some of those fields  
5     may themselves be complex things such as JSON encoded  
6     objects, and the names may not make sense.

7             SPECIAL MASTER GARRIE: Yeah.

8             I think that got at your question; right?

9             MS. WEAVER: Yes. Thank you.

10            MR. EUGENE ZARASHAW: But given -- yes, but we  
11     could find what are the largest tables by size and get  
12     their schemas.

13            SPECIAL MASTER GARRIE: Okay. Any other  
14     questions from plaintiffs?

15            MS. WEAVER: So many, but we'll hold for now.

16            SPECIAL MASTER GARRIE: Okay. Well, just so  
17     you're -- everybody is aware, I see the finish line may  
18     be 20 minutes away, so you best brainstorm away.

19            My next questions are really focused around  
20     the Graph API.

21            Mr. Zarashaw, Mr. Elia, are you guys -- either  
22     of you familiar with the Graph API Explorer application?

23            MR. STEVEN ELIA: I am.

24            MR. EUGENE ZARASHAW: I am.

25            SPECIAL MASTER GARRIE: Okay. How many

## HIGHLY CONFIDENTIAL

Page 149

1 versions of Graph API have been created during the  
2 relevant time period, to your knowledge?

3 MR. STEVEN ELIA: I don't know off the top of  
4 my head, but we typically have three or four different  
5 versions in any calendar year, and we started versioning  
6 in, again, roughly 2015 or so.

7 SPECIAL MASTER GARRIE: And before 2015?

8 MR. STEVEN ELIA: Before then, there weren't  
9 versions to the degree we introduced changes that were  
10 not backwards compatible to APIs. We had different  
11 mechanisms for doing so.

12 SPECIAL MASTER GARRIE: Okay. Has GAE -- I'll  
13 call it GAE -- ever accessed data from sources other  
14 than the social graph?

15 MR. STEVEN ELIA: The Graph API Explorer is  
16 only able to access graph APIs, and so --

17 SPECIAL MASTER GARRIE: Since its inception or  
18 as of 2015?

19 MR. STEVEN ELIA: Since inception.

20 SPECIAL MASTER GARRIE: But I said the social  
21 graph.

22 MR. STEVEN ELIA: Could you repeat the full  
23 question?

24 SPECIAL MASTER GARRIE: So has GAE ever  
25 accessed data from sources other than the social graph?

HIGHLY CONFIDENTIAL

Page 150

1 MR. STEVEN ELIA: So the Graph API can call  
2 any arbitrary Graph API. There are Graph APIs that are  
3 unrelated to the social graph.

4 SPECIAL MASTER GARRIE: So what user data  
5 sources -- so then what are those?

6 MR. STEVEN ELIA: Commonly this would be APIs  
7 that have nothing to do with users. There are APIs, for  
8 example, for managing business assets like ads or pages.

9 SPECIAL MASTER GARRIE: Specific to users, so  
10 the GAE just calls Graph APIs, and then by definition,  
11 it talks to just a social graph as it relates to user  
12 data or user inferred data?

13 MR. STEVEN ELIA: Graph APIs, as they relate  
14 to user data deal with just the social graph, I would  
15 say. I think that was what you were asking.

16 SPECIAL MASTER GARRIE: That is correct.  
17 So is the answer yes?

18 MR. STEVEN ELIA: I would answer yes.

19 SPECIAL MASTER GARRIE: So then what user data  
20 sources did each version of the Graph API have access  
21 to?

22 MR. STEVEN ELIA: Sure. By version.

23 And I think we're back at earlier questioning  
24 and difficulties when it comes to identifying things  
25 like which of these 55 systems might have user data.



HIGHLY CONFIDENTIAL

Page 151

1           SPECIAL MASTER GARRIE: Mr. Zarashaw, do you  
2       agree?

3           MR. EUGENE ZARASHAW: Yes. They're going to  
4       be -- there's, again, the probabilistic answer, which is  
5       most of it is going to come from TAO. But then there's  
6       going to be in each case APIs that have dedicated  
7       systems. And depending on the time we added the  
8       specific API or with which version, the other systems  
9       might have changed.

10           So, again, things like the storage of new  
11       seat posts, such as Timeline -- would be in the  
12       Timeline back end, messages would be in messaging  
13       back end. Photos themselves have the dedicated  
14       back end, while the meta data about the photo might be  
15       in the Timeline back end.

16           SPECIAL MASTER GARRIE: My question is: Is  
17       that data -- and I guess we need to get -- I think the  
18       DYI, your engineer or engineering -- engineers will  
19       answer is that duplicative or not of what's in the DYI  
20       file?

21           MR. EUGENE ZARASHAW: Exactly.

22           SPECIAL MASTER GARRIE: Does Facebook store  
23       data about users it's not directly associated with the  
24       Facebook ID but is capable of being associated with a  
25       specific user?

HIGHLY CONFIDENTIAL

Page 152

1 I think the answer to that we've established  
2 is yes.

3 MR. EUGENE ZARASHAW: Yes.

4 SPECIAL MASTER GARRIE: Am I right, the answer  
5 is "yes"?

6 MR. EUGENE ZARASHAW: The answer is yes, the  
7 IP address associated data.

8 SPECIAL MASTER GARRIE: IP, OS, Windows  
9 update, time and length of chats, length of video,  
10 whatever. There's lots of metadata about data.

11 Are there different processes for different  
12 data?

13 MR. EUGENE ZARASHAW: What do you mean by  
14 processes in this case?

15 SPECIAL MASTER GARRIE: I guess what I'm  
16 thinking here is for data that you infer from a user  
17 profile and a user's action, are there processes that --  
18 like how that data is written out?

19 MR. EUGENE ZARASHAW: I would not call it  
20 processes. It would just depend on the implementation  
21 of the specific use case, where the specific use case  
22 might leverage different pieces of infrastructure and  
23 different systems and flow in different ways, depending  
24 on how the engineers implementing decided to do it.

25 SPECIAL MASTER GARRIE: There's no

HIGHLY CONFIDENTIAL

Page 153

1 documentation here.

2 MR. EUGENE ZARASHAW: There might be some  
3 documentation for specific use cases as they were built.  
4 And generally that documentation falls out of date  
5 rapidly versus the actual evolving systems.

6 SPECIAL MASTER GARRIE: Is there a database or  
7 a central -- or a set of central databases for use  
8 cases?

9 MR. EUGENE ZARASHAW: I don't know of one.  
10 Steven?

11 MR. STEVEN ELIA: Same. I'm not aware of any.

12 SPECIAL MASTER GARRIE: That's a question  
13 worth following up on.

14 So then when I look at this list,  
15 Mr. Zarashaw, of 55 systems, are there any of these  
16 55 systems you can shed light on that we haven't  
17 discussed?

18 MR. EUGENE ZARASHAW: I'm pulling up the list  
19 right now.

20 SPECIAL MASTER GARRIE: I can start at the  
21 top.

22 Ad Market DB?

23 MR. EUGENE ZARASHAW: Last time I looked at  
24 it, it stores data that is advertiser based rather than  
25 user based. So data about advertiser accounts, the ads

HIGHLY CONFIDENTIAL

Page 154

1     they're running, their budgets, and so on.

2             SPECIAL MASTER GARRIE: I did too, which is  
3     why maybe -- how could it -- well, yes. All right.

4             No. 2, Ads Raw Storage. Again --

5             MR. EUGENE ZARASHAW: I don't know what that  
6     one is.

7             SPECIAL MASTER GARRIE: And Mr. Elia, I assume  
8     you don't know as well?

9             MR. STEVEN ELIA: That's right. I don't know.

10            SPECIAL MASTER GARRIE: Archival?

11            MR. EUGENE ZARASHAW: Don't know.

12            MR. STEVEN ELIA: Me neither.

13            SPECIAL MASTER GARRIE: Callisto -- is  
14     Callisto and Calypso different?

15            MR. EUGENE ZARASHAW: I believe so. One of  
16     these is the messaging back end. I don't know what the  
17     other one is.

18            SPECIAL MASTER GARRIE: Got it.

19            And, Mr. Elia, would you by chance know the  
20     difference between Callisto and Calypso?

21            MR. STEVEN ELIA: I don't.

22            SPECIAL MASTER GARRIE: Okay.

23            MR. EUGENE ZARASHAW: Calypso is the messaging  
24     back end. I don't know what Callisto is.

25            SPECIAL MASTER GARRIE: And they would have --

HIGHLY CONFIDENTIAL

Page 155

1 the messaging back end would have its own, like -- just  
2 so I understand, theoretically TAO connect the user --  
3 how does TAO connect to Calypso, or does it?

4 MR. EUGENE ZARASHAW: I don't know that it  
5 even does. It might be a completely separate storage  
6 system or it might be an abstraction layer on top of  
7 another storage system. We would need to dig in to  
8 understand how it's built.

9 SPECIAL MASTER GARRIE: Dumbo?

10 MR. EUGENE ZARASHAW: I don't know.

11 SPECIAL MASTER GARRIE: Not at Disneyland.

12 MR. STEVEN ELIA: I'm not familiar with it  
13 either.

14 SPECIAL MASTER GARRIE: And if counsel for  
15 Facebook, if anybody from Facebook has knowledge of what  
16 any of these systems are, feel free to raise your hand  
17 and Zoom, and I will gladly welcome any insight.

18 ElasticSearch?

19 MR. EUGENE ZARASHAW: It is a text search  
20 service, but I don't know what it's actually used for,  
21 given we also have Unicorn, which can also do this.

22 MR. STEVEN ELIA: And I'm not familiar with  
23 it.

24 SPECIAL MASTER GARRIE: EverStore?

25 MR. EUGENE ZARASHAW: It's a blob store for

HIGHLY CONFIDENTIAL

Page 156

1 storing images and videos and other large blobs.

2 SPECIAL MASTER GARRIE: You have several  
3 patents using that as well actually for that storage.

4 Laser?

5 MR. EUGENE ZARASHAW: It is caching layer in  
6 front of slower forms of storage.

7 SPECIAL MASTER GARRIE: And LogDevice?

8 MR. EUGENE ZARASHAW: Don't know.

9 Steven?

10 MR. STEVEN ELIA: Me neither.

11 SPECIAL MASTER GARRIE: Manifold or Manifold?

12 MR. EUGENE ZARASHAW: Manifold is another blob  
13 storage service. I don't remember if it's -- came  
14 before or after EverStore.

15 SPECIAL MASTER GARRIE: Do we know what's  
16 stored in that one?

17 MR. EUGENE ZARASHAW: If I remember correctly,  
18 photos and videos. I actually don't recall which one is  
19 being transitioned to which.

20 SPECIAL MASTER GARRIE: Again, knowing which  
21 ones transition and which ones don't will be helpful.  
22 Which ones --

23 MR. EUGENE ZARASHAW: It might also just be a  
24 slightly different blob storage case and both of them  
25 are long term, going to be there.

HIGHLY CONFIDENTIAL

Page 157

1 I don't recall the distinction but both do  
2 store blobs. It's akin to Amazon S3.

3 SPECIAL MASTER GARRIE: Okay. All right.  
4 Memcache, we know, we discussed.

5 OIL, we don't know.

6 Omnistore?

7 MR. EUGENE ZARASHAW: No.

8 SPECIAL MASTER GARRIE: Mr. Steven?

9 MR. STEVEN ELIA: I'm not familiar, no.

10 SPECIAL MASTER GARRIE: Poseidon.

11 MR. EUGENE ZARASHAW: No.

12 MR. STEVEN ELIA: Me neither.

13 SPECIAL MASTER GARRIE: RSC Storage?

14 MR. EUGENE ZARASHAW: Nope.

15 MR. STEVEN ELIA: Me neither.

16 SPECIAL MASTER GARRIE: Synapse.

17 MR. EUGENE ZARASHAW: No.

18 MR. STEVEN ELIA: Me neither.

19 SPECIAL MASTER GARRIE: Tableau Server?

20 MR. EUGENE ZARASHAW: I know what it is but  
21 not what we use it for.

22 SPECIAL MASTER GARRIE: I mean, it's a data  
23 analytics tool at its core.

24 MR. EUGENE ZARASHAW: Exactly. I've seen us  
25 use it for things like internal recruiting data and

HIGHLY CONFIDENTIAL

Page 158

1 other IT -- more like IT things rather than production  
2 uses.

3 SPECIAL MASTER GARRIE: Yeah. I'm not exactly  
4 sure where user data fits in there, but...

5 Tally?

6 MR. EUGENE ZARASHAW: It's a counter service  
7 used by the integrity teams to count how many times  
8 various things have happened.

9 SPECIAL MASTER GARRIE: Does it record API  
10 requests?

11 MR. EUGENE ZARASHAW: It might be used to [REDACTED]  
[REDACTED], but it wouldn't be in a particularly  
13 identifiable way.

14 SPECIAL MASTER GARRIE: TAO we know.

15 TimelineDB we know.

16 TW Storage?

17 MR. EUGENE ZARASHAW: Don't know.

18 MR. STEVEN ELIA: I'm not aware.

19 SPECIAL MASTER GARRIE: Okay. UDB?

20 MR. EUGENE ZARASHAW: User databases, so a  
21 sharded set of MySQL databases, sharded by user ID.

22 SPECIAL MASTER GARRIE: So that would know if  
23 we had user data in there.

24 MR. EUGENE ZARASHAW: I don't recall the exact  
25 interplay between UDBs and TAO backing store and which



HIGHLY CONFIDENTIAL

Page 159

1 part is where. [REDACTED]

3 SPECIAL MASTER GARRIE: Before we get to the  
4 remainder of this list, this is just for Facebook.

5 For the other systems that you removed into  
6 these five buckets of internal test systems, systems  
7 Facebook teams, systems that do not serve, systems that  
8 do not store, and other systems, I need you to put of  
9 the -- from the 149 systems which you cut it down, which  
10 ones fall into which bucket; right?

11 So this one falls in this bucket, this one is  
12 in this bucket, just so I have an understanding of how  
13 it was organized.

14 MR. FALCONER: Yeah. Absolutely. Yeah.

15 So, I mean -- yeah. We're more than happy to  
16 do that.

17 And then, you know, the questions on these  
18 sources, you know, we had brought Mr. Elia and  
19 Mr. Zarashaw -- those are the folks we identified as  
20 most knowledgeable and kind of most likely to be able to  
21 answer the Special Master's questions today based on  
22 what was in the orders.

23 SPECIAL MASTER GARRIE: I get it.

24 MR. FALCONER: Anything on these sources, you  
25 know, we veered into a little DYI and advertising and

## HIGHLY CONFIDENTIAL

Page 160

1 privacy, anything on that, we're happy to identify the  
2 right people, get them in front of you so you can ask  
3 the questions you have. Anything like that you need,  
4 we're more than happy to do it.

5 SPECIAL MASTER GARRIE: Yeah. I get it. I  
6 think it was productive. We only have another  
7 25 systems to quickly run through.

8 Up2x?

9 MR. EUGENE ZARASHAW: No.

10 MR. STEVEN ELIA: I'm not familiar.

11 SPECIAL MASTER GARRIE: Warm Storage?

12 MR. EUGENE ZARASHAW: It's an offline backup  
13 system or partly offline backup system.

14 SPECIAL MASTER GARRIE: All right. Well, it  
15 will be very helpful for me from the Facebook lawyers,  
16 remove all the backup systems, all the offline memory  
17 caching, caching systems from this list, because I think  
18 it will get substantially smaller.

19 MR. FALCONER: Right.

20 SPECIAL MASTER GARRIE: XDB we talked about.

21 ZippyDB or -- how do you say that? Akkio  
22 or Akkio?

23 MR. EUGENE ZARASHAW: I don't know what Akkio  
24 is.

25 MR. STEVEN ELIA: I have a high-level

HIGHLY CONFIDENTIAL

Page 161

1 understanding of ZippyDB, although it's not something  
2 I've used. I think it has data for some ephemeral use  
3 cases, things like notifications, for example.

4 SPECIAL MASTER GARRIE: Like notifications to  
5 users, like hurricane is coming?

6 MR. STEVEN ELIA: I'm not sure what kind of  
7 notifications.

8 SPECIAL MASTER GARRIE: But it's outside  
9 messaging. It's push notifications to the user, not  
10 messaging notifications?

11 MR. STEVEN ELIA: Again, I'm not sure which  
12 specific kinds.

13 SPECIAL MASTER GARRIE: Okay.

14 Assistant User Memory? Is that not a caching  
15 tool?

16 MR. EUGENE ZARASHAW: I don't know what that  
17 is. I'm learning a lot about our systems from this  
18 experience.

19 SPECIAL MASTER GARRIE: As am I.

20 Mr. Elia?

21 MR. STEVEN ELIA: I'm not familiar with it  
22 either.

23 SPECIAL MASTER GARRIE: All right. Ibu and  
24 Ibubiz?

25 MR. EUGENE ZARASHAW: I don't know.

HIGHLY CONFIDENTIAL

Page 162

1 MR. STEVEN ELIA: Me neither.

2 SPECIAL MASTER GARRIE: But they're MySQL

3 so...

4 Does counsel know any idea why they're  
5 identified as MySQL? Are they systems within MySQL? Or  
6 are they MySQL systems?

7 Because [REDACTED]. And  
8 like is Ibu and Ibubiz the same or are they unique?

9 MR. EUGENE ZARASHAW: My best guess would be  
10 that [REDACTED]. It's a pretty  
11 involved layer in addition to it. And I don't think it  
12 uses MySQL anymore underneath.

13 On the other hand, we do run some other MySQL  
14 clusters for different purposes, and while the software  
15 may be about the same, the purposes are different. It's  
16 machines with different data on them.

17 SPECIAL MASTER GARRIE: Okay.

18 MR. FALCONER: That's a far better answer than  
19 you would have gotten from counsel, I just want to note  
20 for the record.

21 SPECIAL MASTER GARRIE: Yeah. Noted.

22 TacoSRS.

23 MR. EUGENE ZARASHAW: I don't know.

24 MR. STEVEN ELIA: That may be -- I don't know  
25 if Taco, which is a system renamed itself, in which case

HIGHLY CONFIDENTIAL

Page 163

1 you also might be better able to speak to that than me,  
2 Eugene.

3 MR. EUGENE ZARASHAW: I honestly don't  
4 remember what Taco is anymore.

5 MR. STEVEN ELIA: It's very similar to TAO but  
6 I think had fewer reliability guarantees perhaps.

7 MR. EUGENE ZARASHAW: Oh, is that -- oh, I do  
8 recall what Taco is now. That was the Memcache base TAO  
9 without the MySQL part. So data went in and it might  
10 just disappear at any time. It was meant for more  
11 convenient caching than Memcache directly.

12 SPECIAL MASTER GARRIE: All right. So we can  
13 junk that one.

14 All right. Ads Online Storage?

15 MR. EUGENE ZARASHAW: Don't know.

16 MR. STEVEN ELIA: Me neither.

17 SPECIAL MASTER GARRIE: I don't even know how  
18 to say the next one.

19 MR. EUGENE ZARASHAW: It's okay. It's a no  
20 for the next three as well.

21 MR. STEVEN ELIA: Same.

22 SPECIAL MASTER GARRIE: Yeah. I don't even  
23 know.

24 Orderdb?

25 MR. EUGENE ZARASHAW: Don't know what that one

HIGHLY CONFIDENTIAL

Page 164

1 is.

2 The next one I know anything about is  
3 Payments.

4 MR. STEVEN ELIA: Yeah. I don't know the ones  
5 before Payments.

6 SPECIAL MASTER GARRIE: Payments.

7 MR. EUGENE ZARASHAW: Any time we bill a user  
8 for something, whether we bill them for ads, for  
9 example -- the most common cases, we charge them for  
10 running ads. But also bill them for sending money to  
11 their friend.

12 SPECIAL MASTER GARRIE: And then you also have  
13 like -- they can buy photos or whatever else services.

14 Is that all run through Payments?

15 MR. EUGENE ZARASHAW: I don't know if you can  
16 buy photos anymore. But, for example, there's  
17 remittances, which you can send money to a friend and  
18 that runs through Payments.

19 SPECIAL MASTER GARRIE: All right. Anything  
20 you want to add there, Steven?

21 MR. STEVEN ELIA: Nope.

22 SPECIAL MASTER GARRIE: Okay.  
23 Presence Storage Services?

24 MR. EUGENE ZARASHAW: I don't know anything up  
25 until 44.

## HIGHLY CONFIDENTIAL

Page 165

1 SPECIAL MASTER GARRIE: Okay. I'll take that.

2 How about you, Steven?

3 MR. STEVEN ELIA: Same.

4 SPECIAL MASTER GARRIE: MySQL.

5 MR. EUGENE ZARASHAW: Likely on there because  
6 we probably have a set of MySQL machines somewhere that  
7 are not covered by the other MySQL clusters listed.

8 SPECIAL MASTER GARRIE: That's what I'm a  
9 little confused about. There are a bunch of MySQL  
10 clusters here that seem -- so one of the things -- these  
11 are systems that plaintiffs identified as storing user  
12 data. But wouldn't MySQL be the like -- MySQL is --  
13 it's a database; right? It's not a --

14 MR. EUGENE ZARASHAW: Right. And it's hard to  
15 tell which instance of that database this is and what  
16 would be in it.

17 MR. STEVEN ELIA: I would add that as Eugene  
18 has previously said, these storage systems are often  
19 implemented and organized here by use case. And so  
20 those other MySQL instances listed here have their own  
21 very specific use cases.

22 SPECIAL MASTER GARRIE: Where would we find  
23 these use cases for these systems?

24 MR. EUGENE ZARASHAW: I recommend going the  
25 other way and looking at what systems are actually big

## HIGHLY CONFIDENTIAL

Page 166

1 enough to serve a couple of -- two billion people.  
2 Because if you're looking for user scale systems that  
3 are storing that much data and having that much load, a  
4 cluster of 50 machines is just not going to do it.

5 So getting some idea of the number of machines  
6 involved might help you sift through these --

7 SPECIAL MASTER GARRIE: Okay. That could be  
8 very helpful. That is actually an excellent idea.

9 We'll give that some more thought.

10 Anything to 46, either one of you?

11 MR. EUGENE ZARASHAW: No. I'm lost until 50.

12 MR. STEVEN ELIA: Same.

13 SPECIAL MASTER GARRIE: We discussed Unicorn.  
14 Research Home is something of interest to me.  
15 What is Research Home?

16 MR. EUGENE ZARASHAW: I don't know.

17 MR. STEVEN ELIA: I don't either.

18 SPECIAL MASTER GARRIE: Does counsel for  
19 Facebook have any idea why Research Home would be on  
20 this list?

21 MR. FALCONER: I don't want to speak off the  
22 cuff on that, no. I mean, we can find out. We're happy  
23 to find out as well, but I don't want -- you know,  
24 if were to say anything inaccurate.

25 SPECIAL MASTER GARRIE: No. I got it.



## HIGHLY CONFIDENTIAL

Page 167

1           The reason I ask is because there's been  
2       testimony from -- that there is no central -- or  
3       testimony and statements from Facebook that there's no  
4       central research repository.

5           When I read the words Research Home, it  
6       strikes me as that may be a central research repository.  
7       I could be wrong.

8           MR. EUGENE ZARASHAW: My guess on that one --  
9       and, again, this would only be a guess -- is this is a  
10      database backing the research dot Facebook dot com site  
11      containing the papers that we publish.

12          SPECIAL MASTER GARRIE: That's what I thought  
13      likely as well.

14          MS. RING: Just to be clear, we don't know.  
15      I mean, you asked, Special Master, counsel. There are  
16      crazy names for things that make no sense and end up  
17      having no relationship to what's actually in a data  
18      source. I've seen that for over a decade. So we will  
19      find out. But it would not be out of the ordinary for  
20      it to have nothing at all to do with research.

21          SPECIAL MASTER GARRIE: I mean Tacos -- don't  
22      have to do with Tacos.

23          MS. RING: Exactly. Or Dumbo. I mean --

24          SPECIAL MASTER GARRIE: Unicorns. My daughter  
25      that was very excited when I told her I was going to

HIGHLY CONFIDENTIAL

Page 168

1 talk about unicorns today.

2 So that --

3 MR. EUGENE ZARASHAW: Of all the things I know  
4 on that list.

5 SPECIAL MASTER GARRIE: Anything?

6 MR. STEVEN ELIA: I don't specifically know  
7 52 PYMK Leaf Tailer. But PYMK does stand for "people  
8 you may know." It's the part of the product that  
9 suggest friends to you. And so while I don't know what  
10 data is in there, I would guess that it's in service of  
11 that.

12 SPECIAL MASTER GARRIE: Any others?

13 MR. STEVEN ELIA: That's the last I recognize  
14 anything about.

15 SPECIAL MASTER GARRIE: Got it.

16 And then below, there's a drawing and it was  
17 stated that the social graph connects to 12 specific  
18 systems.

19 One thing I did see that was removed from the  
20 list, just for notification, was Hive.

21 But one thing I had a question about is do  
22 you -- for these four -- for those 12 systems, right,  
23 Ad Market DB -- which I'm not quite sure I understand --  
24 Callisto, EverStore, Laser, Manifold, I think we can  
25 probably get rid of Memcache, MySQL, again --

HIGHLY CONFIDENTIAL

Page 169

1 MR. EUGENE ZARASHAW: Apologies. I need to  
2 update the record.

3 It is Callisto, not Calypso, that is the  
4 messaging store.

5 SPECIAL MASTER GARRIE: Yeah.

6 Tally, TAO.

7 So -- all right. I think this was helpful.  
8 I think we'll have to add one -- I'll get a more  
9 specific set of requests for further clarification, as  
10 well as the other things that I mentioned. I'll try to  
11 get an order out, sort of clarifying that.

12 And I may tier the approach.

13 Does plaintiffs have any further questions  
14 they'd like to ask?

15 MS. WEAVER: I think we're okay right now.

16 Thank you.

17 SPECIAL MASTER GARRIE: Facebook, any  
18 questions you would like to ask?

19 MR. FALCONER: Not right now. No. Thank you  
20 for the opportunity.

21 But I think what we'll do is, you know, once  
22 we get the transcript, we'll go through it and if we  
23 have any clarification that we think it would be helpful  
24 to provide to the Special Master, we'll do that.

25 Probably just with a letter through the JAMS portal, you

HIGHLY CONFIDENTIAL

Page 170

1 know, so everybody gets a copy. But I don't think  
2 anything for this afternoon.

3 SPECIAL MASTER GARRIE: All right. We can go  
4 off the record for one second.

5 (Pause in proceedings.)

6 SPECIAL MASTER GARRIE: All right. We'll go  
7 back on the record.

8 All right. So rather than issuing a formal  
9 order for the additional requested information, I'm  
10 going to post via JAMS access my request for additional  
11 information, and likely to request the scheduling of an  
12 additional hearing. All hearing will be labeled as  
13 confidential or actually highly confidential.

14 Any party that does want to attend this  
15 hearing has to sign the protective order and receive --  
16 and notified accordingly.

17 In addition, I'm requesting that the  
18 plaintiffs, if they come up with additional systems or  
19 any particular systems they think would be of value, to  
20 not let the five days that stand after the issuance of  
21 my order, identify them as early as possible would be  
22 much appreciated.

23 And otherwise I believe that is all. And I  
24 thank the parties for making Mr. Zarashaw and Mr. Elia  
25 available for today's hearing, and it was extremely

HIGHLY CONFIDENTIAL

Page 171

1 helpful and informative. I feel like we were able to  
2 cut down the systems, at least from the list of 55  
3 fairly substantially and develop clear parameters and  
4 focus. So I do appreciate it.

5 And I'm going to move as fast as possible to  
6 get to the bottom -- to resolve this issue.

7 MS. RING: Special Master Garrie, if I can  
8 raise one more thing, and I don't know if we can resolve  
9 it today but at least raise it and then think about how  
10 to resolve it.

11 I think given the nature of the things that  
12 are being discussed, we would like the hearings to be  
13 sealed. Not just to people who are here to sign the  
14 protective order, but to seal. I mean, this is very  
15 detailed information about all of the companies'  
16 internal systems, so --

17 SPECIAL MASTER GARRIE: I'm fully with you.  
18 So on the record, I'm fully in support of that, but that  
19 is a motion Facebook has to bring to the Court for the  
20 Court to issue a ruling.

21 MS. RING: Okay.

22 SPECIAL MASTER GARRIE: Accordingly, I do  
23 agree with you that we are covering the inner workings  
24 of Facebook systems and databases and technology and  
25 systems.

HIGHLY CONFIDENTIAL

Page 172

1 But, again, that's a relief you would -- a  
2 motion you would have to bring part before the judge or  
3 the Court, and according to the local rules, follow the  
4 respective procedures that underline that and the Court  
5 will rule accordingly.

6 MS. RING: Understood. I just wanted to -- I  
7 didn't want to leave the hearing without having  
8 addressed that.

9 SPECIAL MASTER GARRIE: No. And that is why  
10 today's hearing, there was such -- there was a lot of  
11 interest and request to attend.

12 And I would also request the plaintiffs to  
13 make sure and inform your -- and inform your fellow  
14 plaintiffs' lawyers to channel everything through you  
15 pursuant to the existing order.

16 MS. WEAVER: That has just been done.

17 SPECIAL MASTER GARRIE: Well, just to  
18 communicate it accordingly, because now that judge --

19 Well, we'll go off the record now. We're  
20 done.

21 (Whereupon, proceedings adjourned at 5:28 p.m.)

22 ---o0o---

23

24

25

HIGHLY CONFIDENTIAL

Page 173

1 STATE OF CALIFORNIA )

)

2 COUNTY OF YOLO )

3 ---o0o---

4 I, Katy E. Schmidt, a Certified Shorthand  
5 Reporter, do hereby certify:

6 That said proceedings were taken before me at  
7 the time and place therein set forth and were taken down  
8 by me in shorthand and thereafter transcribed into  
9 typewriting under my direction and supervision;

10 I further certify that I am neither counsel  
11 for, nor related to, any party to said proceedings, and  
12 am not in any way interested in the outcome thereof.

13 In witness whereof, I have hereunto subscribed  
14 my name.

15 Dated: February 18, 2022

16

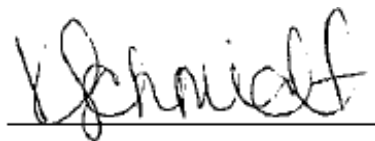
17

18

19

20

21



22

Katy E. Schmidt

RPR, RMR, CRR, CSR 13096

23

24

25

## HIGHLY CONFIDENTIAL

[&amp; - acronyms]

Page 1

<b>&amp;</b>	<b>2019</b> 72:11 137:5	<b>555</b> 3:11	<b>access</b> 10:14,14,21
<b>&amp;</b> 3:9,15	<b>2022</b> 1:14 2:24 4:2	<b>5:28</b> 172:21	10:22,24 11:2,3,6
<b>1</b>	173:15	<b>6</b>	11:9,12,17,17,19
<b>1</b> 5:22 24:10 30:2	<b>2100</b> 3:19	<b>6</b> 42:22	11:23 12:2,7,14,14
30:6 124:18 125:8	<b>214.698.3170</b> 3:20	<b>60</b> 138:21,23	13:3,9,19 15:2,15
<b>1.0</b> 12:16 13:13	<b>25</b> 160:7	<b>64</b> 61:6,16,18,20	15:16 18:10,12,18
<b>10</b> 147:12,20,24,25	<b>3</b>	62:18	19:1 20:20 22:1
<b>12</b> 25:15 30:16	<b>3</b> 6:2,6 24:10	<b>7</b>	32:14 35:12 36:11
42:6 43:22 168:17	125:8	<b>75201</b> 3:19	37:9 48:21 79:15
168:22	<b>30</b> 138:20	<b>9</b>	80:4,22 81:6,10,11
<b>1201</b> 3:7	<b>3200</b> 3:7	<b>9</b> 134:4	81:12,16,20,24,25
<b>128</b> 142:16,17	<b>374</b> 139:11 142:17	<b>95</b> 86:20,23	90:5 92:6 106:8
144:25	<b>38</b> 43:9	<b>98101-3052</b> 3:7	107:5 109:15
<b>12th</b> 3:11	<b>4</b>	<b>994607</b> 3:11	110:15 114:13
<b>13096</b> 1:24 2:25	<b>4</b> 24:10 132:12	<b>9th</b> 7:22	117:9 125:13
173:22	139:11	<b>a</b>	128:7,24 133:23
<b>149</b> 4:11 159:9	<b>40</b> 120:8	<b>abbreviate</b> 100:19	133:25 136:5
<b>15/2016</b> 137:14	<b>415.445.4003</b> 3:12	<b>ability</b> 15:23 86:4	137:11,14,16,21
<b>1600</b> 3:11	<b>42</b> 27:15 92:19	107:18 110:24	140:16 145:6
<b>17</b> 1:14 2:24 4:2	<b>43</b> 25:17,19 27:11	126:13 140:5	149:16 150:20
<b>17973</b> 173:21	42:5,7 77:5 87:8	<b>able</b> 13:11,12 15:5	170:10
<b>18</b> 139:1 173:15	92:9 98:6	18:24 19:1 22:19	<b>accessed</b> 95:10
<b>1:07</b> 2:24	<b>44</b> 164:25	23:17 31:9 40:15	117:14 121:7
<b>2</b>	<b>46</b> 166:10	43:17 47:2 48:3	136:4 149:13,25
<b>2</b> 5:23 24:10 30:2	<b>48</b> 44:1	54:14 78:13 115:3	<b>accessible</b> 10:23
30:6 84:25 85:5	<b>487</b> 144:25	128:13 129:15,20	11:2 21:24 34:5
154:4	<b>495</b> 142:16	130:6 149:16	35:18 51:3
<b>2.0</b> 12:9	<b>5</b>	159:20 163:1	<b>accessing</b> 13:6
<b>20</b> 29:10 86:22	<b>50</b> 27:11 166:4,11	171:1	23:8 81:17
137:9 148:18	<b>500</b> 86:11	<b>absolutely</b> 82:13	<b>accomplish</b> 102:4
<b>2001</b> 3:19	<b>5096673</b> 1:25	159:14	<b>account</b> 79:13
<b>2013</b> 69:24	<b>52</b> 168:7	<b>abstracted</b> 38:7	105:11,14,17
<b>2014</b> 69:25 113:23	<b>55</b> 4:13 7:11,14	<b>abstraction</b> 155:6	<b>accounts</b> 37:9
113:23,24 114:6,8	25:13 49:8 62:24	162:7,10	153:25
137:2	65:23 66:12 98:6	<b>accelerate</b> 134:8	<b>accurate</b> 128:9
<b>2015</b> 26:25 31:9	110:20,23 120:8	<b>accept</b> 109:20,24	129:22
149:6,7,18	131:20 132:12	110:5	<b>acquire</b> 79:4
<b>2016</b> 26:25 31:9	150:25 153:15,16	<b>acceptable</b> 7:24	<b>acquired</b> 32:12
<b>2018</b> 17:12 20:13	171:2	17:5,14	130:22
72:11 137:5			<b>acronyms</b> 100:9



## HIGHLY CONFIDENTIAL

[act - anybody]

Page 2

<b>act</b> 81:13 103:19 <b>acting</b> 13:15 <b>action</b> 152:17 <b>actions</b> 135:7 <b>active</b> 105:14,15 <b>actively</b> 33:11 <b>activities</b> 5:24 6:1 6:8 46:7 62:12,14 74:9 119:22 124:20 <b>activity</b> 5:19,22 6:4,9,9 20:17,23 20:25 24:15,17 25:20 30:14 46:6 46:8 48:8,12 51:13 52:19 53:9 53:10 62:10 69:11 72:16 73:8 74:3 75:5,21,24 78:20 79:5 82:9 119:21 119:23 123:13,22 124:2,13,25 125:2 143:7 144:15 <b>actual</b> 25:22 28:25 56:13 135:18 153:5 <b>ad</b> 20:25 28:10,10 28:21,25 29:1,4,7 29:8,10,12,15,16 29:21,22 30:3,4 31:6,14,15 47:4,19 48:3,4,6,9,22,24 49:4,25,25 51:7,7 65:11 69:23 74:1 75:19 76:2 77:2,7 132:22 142:14 153:22 168:23 <b>add</b> 10:9 35:7 59:20 83:2 87:18 95:11 135:24 137:18,24 164:20	165:17 169:8 <b>added</b> 151:7 <b>adding</b> 63:4 <b>addition</b> 131:14 162:11 170:17 <b>additional</b> 25:3 51:1 60:22 112:15 116:19,20 128:7 170:9,10,12,18 <b>address</b> 24:3 26:1 39:19,20,25 41:17 42:2 43:8 152:7 <b>addressed</b> 172:8 <b>addresses</b> 23:7 39:17 41:9,12 42:12 51:6 65:9 76:14,18 <b>adjourned</b> 172:21 <b>ads</b> 28:9,18,20 31:7 47:12 50:12 50:13 56:13 57:10 67:12 70:10 76:5 76:8 78:8,8,23 123:9,16 133:1,4,5 133:8,24 142:7 150:8 153:25 154:4 163:14 164:8,10 <b>advertisements</b> 133:6 <b>advertiser</b> 29:22 47:19 153:24,25 <b>advertising</b> 133:7 142:8 159:25 <b>advocacy</b> 7:9 <b>affirmed</b> 49:19 <b>afternoon</b> 170:2 <b>age</b> 65:6 <b>aggregate</b> 86:23 89:12,12	<b>aggregated</b> 4:22 78:19 88:9 89:1 89:13 92:13 <b>agnostic</b> 87:17 <b>ago</b> 61:17 107:17 <b>agree</b> 39:8 48:17 49:14 131:2,3 151:2 171:23 <b>agreement</b> 15:14 16:8 113:18 <b>agreements</b> 125:21 126:1 <b>ahead</b> 8:6 40:17 57:14 66:2 81:9 113:3 133:9 137:12 138:11 <b>akin</b> 157:2 <b>akkio</b> 160:21,22 160:23 <b>alive</b> 32:9 <b>allow</b> 11:18,22 13:5 19:11 62:17 97:17 110:16 128:17,22 <b>allowed</b> 14:21 101:12 110:1 125:24 139:23 <b>allowing</b> 112:23 137:16 <b>allows</b> 13:2 19:15 44:19 106:7,23 107:6 <b>alternative</b> 102:25 <b>amazon</b> 74:11 75:7 79:6 82:7,9 82:12,13 125:15 157:2 <b>amount</b> 40:14 <b>analogy</b> 40:16,18 41:8 121:20	<b>analysis</b> 27:2,4 29:18 <b>analytics</b> 27:6 157:23 <b>analyze</b> 84:16 <b>android</b> 74:15 <b>angeles</b> 4:1 <b>animal</b> 21:18 <b>anne</b> 3:10 <b>anonymize</b> 105:8 <b>anonymized</b> 4:23 4:24 23:9 41:22 61:8 76:4,4 98:16 98:18,19 105:11 <b>anonymizing</b> 103:12 <b>answer</b> 8:25 12:23 15:25 17:15 20:10 20:15 23:12 24:7 32:19 34:6,23,25 35:3 42:16 46:11 46:19,21,25 47:2 50:4,18 52:9 54:1 59:9 65:15 68:15 77:9 85:12,12 86:19 103:24 104:9 117:21 129:15 130:16 132:10 143:25 146:18 150:17,18 151:4,19 152:1,4,6 159:21 162:18 <b>answered</b> 7:22 46:16,17 <b>answering</b> 8:21 53:18 <b>answers</b> 7:25 70:3 102:16 104:11 <b>anybody</b> 17:1 77:14 107:22 155:15
--	--	--	--

## HIGHLY CONFIDENTIAL

[anybody's - assumption]

Page 3

<b>anybody's</b> 58:5 <b>anymore</b> 35:12 72:1 162:12 163:4 164:16 <b>anyways</b> 117:16 139:3 141:17 <b>api</b> 8:10 9:5,15,18 9:21,23 10:6,14,14 12:8,22 17:24 18:1,3,5,8,15 19:5 19:22,23 74:6,6,8 74:12,13,15 75:5,7 75:11 76:2 79:16 79:23,25 80:2,5,10 80:24 82:2,24 83:15,16 84:19,20 85:10,19 86:4 110:13,16 111:9 111:10,11 112:8,8 112:12 113:6,7,12 113:13,15 114:3 114:17,19,20 115:1 116:10,15 116:15,21 117:6,7 117:24 118:7 122:14,15,19,20 125:7 126:15,24 127:12,13,15,19 128:11,14,15,25 129:1,13,23 130:2 130:25 135:5,7,13 135:13 136:6,8,9 136:10,12,14,15 136:18,25 137:1,8 137:9,10,10 138:2 138:2 140:6 148:20,22 149:1 149:15 150:1,2,20 151:8 158:9,12 <b>api's</b> 9:13 84:23 114:24	<b>apis</b> 9:11,17 10:10 10:11 11:22 13:4 18:2 20:1 79:19 79:22,24 80:24 81:20 84:20 85:7 85:19 86:2,7,9,11 86:14 106:5,8 110:16,18 113:2,4 114:12 115:7,10 115:25 117:4,14 118:6 122:7 123:4 124:5,7,21,22,23 126:6 127:1 128:2 129:11,11,19 130:25 131:5,6,9 136:18 137:15,19 137:19 138:3 139:12,23 140:23 140:25 149:10,16 150:2,6,7,10,13 151:6 <b>apocalypse</b> 135:21 <b>apologies</b> 169:1 <b>apologize</b> 16:13,18 16:21 22:9 55:5 58:4 101:3 124:3 <b>app</b> 9:23,24 10:1 10:20,21,22,24 11:19 12:10,10 13:9,11,12,13,15 13:16,18 14:9,13 14:16,17 15:1,15 18:17,19,25 19:8 19:12,13,14 21:4 23:16 80:23,25 81:6,11,14,20,24 81:25 107:4 108:16,21 110:1 111:15,20 112:14 112:17 113:8 114:8 135:14	137:20 143:22 <b>appeal</b> 36:21 37:3 <b>appeals</b> 37:10 <b>appear</b> 53:12 95:17 117:8 <b>appearances</b> 3:1 <b>appeared</b> 3:3,14 <b>appears</b> 89:21 <b>application</b> 11:9 11:10,11,12,13 13:25 14:3,9,18 17:22 26:6 38:4 44:20 51:9 97:25 106:23 107:1,8 111:23 148:22 <b>applications</b> 9:14 106:8 <b>applied</b> 105:25 <b>appreciate</b> 7:8 33:25 50:2 92:5 171:4 <b>appreciated</b> 170:22 <b>approach</b> 169:12 <b>approaches</b> 28:23 <b>appropriate</b> 4:11 9:3,3 <b>approximately</b> 120:20 <b>apps</b> 12:13 14:22 15:4 81:17 96:14 108:10,11 112:21 144:3 <b>arbitrary</b> 147:6 150:2 <b>archival</b> 154:10 <b>archive</b> 40:19,24 40:25 <b>area</b> 8:12 47:25 50:15 59:9 104:5 134:18 140:20	<b>argument</b> 50:5 <b>arrive</b> 71:22 <b>articulate</b> 49:23 71:16 <b>artifact</b> 89:7,7 <b>artifacts</b> 83:23 87:12,15 89:9 <b>asked</b> 4:8 8:24 55:19 167:15 <b>asking</b> 145:2,3 150:15 <b>aspects</b> 139:25 <b>assets</b> 6:12 150:8 <b>assistant</b> 161:14 <b>assisting</b> 7:2 <b>associate</b> 26:13 65:18 77:3 93:23 105:9 120:2 <b>associated</b> 23:21 29:9 31:16 39:4 39:14,15 40:6,6 47:16,18,19 48:12 57:3 59:16,22 62:17 64:15,24 65:2 69:3 73:25 88:23 98:10 103:6 105:21 151:23,24 152:7 <b>associating</b> 105:12 105:13 <b>assume</b> 6:17 26:5 26:22 32:25 36:4 43:7 58:9 78:12 81:24 90:11 94:8 117:15 118:7 129:20 131:25 154:7 <b>assuming</b> 59:9 89:15 <b>assumption</b> 118:12
--	--	---	--

## HIGHLY CONFIDENTIAL

[attached - brought]

Page 4

<b>attached</b> 48:8 75:8 <b>attempts</b> 34:21 <b>attend</b> 170:14 172:11 <b>attorney</b> 6:21 7:3 <b>audience</b> 76:1,21 <b>auld</b> 3:9 <b>auth</b> 112:13 <b>authenticate</b> 80:5 <b>authentication</b> 81:5 <b>author's</b> 81:7 <b>authorization</b> 81:5 <b>authorized</b> 81:12 107:4 <b>auto</b> 109:19,24 <b>automatically</b> 115:3,6,22 <b>available</b> 57:8 74:12 75:6,7 80:1 106:9 113:7 117:14 125:23 138:1,2 144:15 170:25 <b>avenue</b> 3:7,19 <b>aware</b> 35:5 72:22 73:18 75:12 76:25 79:18 87:24 123:4 123:16 126:13,15 126:23 129:8 137:25 148:17 153:11 158:18	59:13 63:21 66:4 67:7 74:7 83:7 89:6 97:7 98:22 101:14,21 107:19 117:20 126:22 131:7,9,11,14,17 131:20,24 132:8 132:11 134:21 135:1,17,19 141:15 146:5,9 150:23 151:12,13 151:14,15 154:16 154:24 155:1 170:7 <b>backed</b> 38:17 130:25 131:6,10 131:11 <b>backing</b> 38:20 121:9 158:25 167:10 <b>backup</b> 25:6 160:12,13,16 <b>backwards</b> 126:3 149:10 <b>bank</b> 32:12 <b>banking</b> 32:6,25 67:17 68:16 <b>banks</b> 32:7 <b>base</b> 86:2 163:8 <b>based</b> 11:22 12:22 13:4 17:16 26:5 31:21 39:17 50:24 56:2 66:9 67:15 81:6,7,11 88:20 89:17 90:22 91:3 111:14 121:3 125:4 136:18 153:24,25 159:21 <b>basic</b> 46:23 71:11 <b>basically</b> 13:2 62:2 75:22 86:25	112:19 113:17 117:1 147:5 <b>batch</b> 75:16 <b>bathroom</b> 55:4 <b>batted</b> 10:6 <b>bearer</b> 81:4,12 <b>beginning</b> 80:6 <b>behalf</b> 10:1 11:13 13:15 81:13 <b>behavior</b> 35:12 72:20 107:2 112:8 113:8 115:12 118:18,21 123:13 124:15 127:13 <b>behaviors</b> 115:11 <b>believe</b> 17:11 20:24 21:15 24:25 43:21 46:24 54:14 66:13 70:10 91:7 98:2 104:17 107:16 108:10 109:25 135:1,8 137:15 144:7 154:15 170:23 <b>benefit</b> 8:8 87:21 <b>benjamin</b> 3:6 <b>best</b> 20:15 43:22 52:22 104:11 148:18 162:9 <b>better</b> 13:7 17:14 38:11 45:11 50:17 63:8 72:25 129:15 142:10 162:18 163:1 <b>beyond</b> 35:8 79:7 134:13 <b>bfalaw.com</b> 3:12 <b>bidirectional</b> 127:1 <b>big</b> 30:18 165:25	<b>bill</b> 164:7,8,10 <b>billion</b> 166:1 <b>billions</b> 30:22 <b>binary</b> 94:10 <b>bing</b> 126:5,7,10,11 126:14,22,24 <b>bit</b> 15:25 33:21 42:15 50:16 54:23 61:6,16,18,20 62:18 130:1 <b>bleichmar</b> 3:9 <b>blob</b> 94:11 155:25 156:12,24 <b>blobs</b> 63:6 147:7,8 156:1 157:2 <b>bluetooth</b> 21:22 120:1 144:23 <b>bodies</b> 128:2 <b>body</b> 116:12 <b>book</b> 24:1 <b>bottom</b> 171:6 <b>box</b> 122:2 <b>brainstorm</b> 148:18 <b>break</b> 26:11 55:1,7 55:10,14,24 100:22,23 101:1,5 101:8 125:11 145:11,13,21,25 146:3 <b>breaking</b> 55:6,16 100:21 <b>breakout</b> 146:1 <b>breathing</b> 52:3 <b>briefly</b> 72:6 <b>bring</b> 50:7 171:19 172:2 <b>broader</b> 32:17 46:9 69:18 <b>brought</b> 159:18
<b>b</b>			
<b>b</b> 53:11,13 <b>b's</b> 53:12 <b>back</b> 7:25 28:5 31:9 33:8 34:25 35:19 36:3 42:3 44:8,17 50:14,17 55:10 57:17,21 58:5 59:4,6,10,12			

## HIGHLY CONFIDENTIAL

[brown - change]

Page 5

<b>brown</b> 88:17,22	<b>cached</b> 112:4	<b>capability</b> 37:21	136:5 153:3,8
<b>browser</b> 9:25 15:8	<b>caching</b> 95:8,22	37:23 63:24 96:13	161:3 164:9
26:2 71:23 72:20	99:7 134:12 156:5	97:19 106:23	165:21,23
108:12 109:17	160:17,17 161:14	107:6 110:1,9,14	<b>categories</b> 56:20
<b>browsers</b> 80:8	163:11	111:4,5,7,8,12,15	<b>caught</b> 65:15
<b>bucket</b> 106:4	<b>calendar</b> 149:5	111:19,24 112:9	<b>cause</b> 26:22 46:20
125:8 159:10,11	<b>california</b> 3:11 4:1	112:10,14 113:5,9	47:5 74:2 97:11
159:12	103:6,19 173:1	113:11,14,14,17	125:14
<b>buckets</b> 125:8	<b>call</b> 11:10 12:8	113:22 114:3,7	<b>caused</b> 36:23
159:6	37:22,23 39:5	115:12 127:17	<b>ccpa</b> 103:15
<b>budgets</b> 154:1	52:11,15 60:5	<b>capable</b> 39:14	<b>centers</b> 38:14 96:2
<b>build</b> 14:13 60:2	75:8 86:2,5 98:6	151:24	<b>central</b> 153:7,7
64:5 77:12 86:4	111:2 112:8,8,12	<b>caps</b> 114:13	167:2,4,6
91:18 108:10	117:7 137:23	<b>capture</b> 21:12,15	<b>centralized</b> 60:18
112:24 138:9	149:13 150:1	29:1	134:18
<b>building</b> 32:7	152:19	<b>captured</b> 29:8	<b>certain</b> 23:15
112:21,23 125:21	<b>called</b> 10:11 14:22	123:14	39:23 46:1 50:23
136:25	32:13 39:6 43:10	<b>capturing</b> 74:4	60:19 63:1,25
<b>builds</b> 14:10	74:21 83:16 102:9	<b>cardinality</b> 31:6	85:12 89:16 90:11
<b>built</b> 38:14 60:20	107:20 111:3,11	<b>care</b> 61:23 98:18	90:24 96:14 99:14
63:1,14,23 77:18	113:21 115:11	<b>cari</b> 3:5	108:2 115:12
86:11 87:19	117:24 121:10,11	<b>carlos</b> 57:20,23	131:5
118:16 138:5	124:7 127:20	101:13	<b>certainly</b> 8:17,19
139:20 153:3	129:18 147:3	<b>carry</b> 127:7	55:6 66:1 72:19
155:8	<b>calling</b> 52:12 84:8	<b>case</b> 7:9,23 16:25	79:12 85:11 88:2
<b>bulk</b> 9:15,16,21	84:23 86:8 117:24	29:15 37:2 40:12	<b>certainty</b> 39:21
10:6 27:1 75:14	<b>callisto</b> 132:12	51:19 54:6 63:5	<b>certified</b> 173:4
75:16	154:13,14,20,24	63:11 64:5 68:13	<b>certify</b> 173:5,10
<b>bunch</b> 94:17	168:24 169:3	77:24 87:20,21	<b>cetera</b> 21:1 125:16
144:24 145:1	<b>calls</b> 9:23,24 83:15	88:15 91:7 99:25	<b>challenge</b> 47:23
165:9	129:2 130:3	109:13 125:20	53:17 54:13 60:12
<b>business</b> 50:5,12	150:10	127:18 128:21	62:20 77:11
142:7,8,12 150:8	<b>calypso</b> 140:9	140:7 151:6	127:25 144:16
<b>button</b> 69:13 72:7	154:14,20,23	152:14,21,21	<b>challenges</b> 64:2
88:24	155:3 169:3	156:24 162:25	<b>challenging</b> 46:18
<b>buy</b> 164:13,16	<b>campaign</b> 76:21	165:19	124:11
<b>bypasses</b> 18:23	<b>capabilities</b> 97:21	<b>cases</b> 15:6 63:3	<b>chance</b> 23:12
<b>c</b>	97:24 109:12	66:5 68:11 77:19	31:24 104:24
<b>c</b> 73:23	111:6 112:20,22	77:23 87:25 99:23	154:19
<b>cache</b> 44:1 134:6,9	114:5,7 127:9	123:15 127:14	<b>change</b> 54:21
	137:16	129:21 130:6,9,17	

## HIGHLY CONFIDENTIAL

[changed - connect]

Page 6

<b>changed</b> 151:9	142:25 167:14	123:22 124:1,12	<b>compile</b> 45:23,25
<b>changes</b> 12:9	171:3	142:22 143:1,6	<b>complete</b> 54:4
149:9	<b>cleared</b> 50:7	144:3,8,14,18,24	81:23 119:15
<b>channel</b> 172:14	<b>clearest</b> 27:25	<b>collection</b> 19:24	<b>completed</b> 138:23
<b>characteristics</b>	<b>clearly</b> 49:20	20:18,23 52:8	<b>completely</b> 17:5
63:2 66:9 90:18	52:13 119:24	<b>collects</b> 14:16	17:14 63:19
90:24 91:4 93:18	143:5	20:24	143:16 155:5
<b>charge</b> 164:9	<b>click</b> 21:11,13,14	<b>column</b> 147:3,8	<b>complex</b> 28:20
<b>chat</b> 58:9 68:8	73:9	<b>com</b> 82:18 167:10	38:4 148:5
93:14,15 107:14	<b>clicked</b> 51:14	<b>come</b> 37:19 55:10	<b>compliant</b> 121:15
107:15,19,19	<b>clicks</b> 20:25 21:8	69:14 86:18 151:5	<b>complicated</b> 46:22
108:2,3 125:5	21:10 28:25 110:5	170:18	49:23 71:14
<b>chats</b> 152:9	<b>client</b> 6:21 7:3	<b>comes</b> 58:23 73:2	<b>comprehensive</b>
<b>cheaper</b> 32:10	10:12 107:16	73:21 85:20 109:6	117:25
95:21	112:24	142:1 150:24	<b>compute</b> 128:16
<b>check</b> 26:6,12	<b>clients</b> 32:8 108:2	<b>comfortable</b> 8:21	<b>computer</b> 22:15
73:20 82:23	<b>close</b> 119:14	<b>coming</b> 37:14 74:4	25:25 26:11 52:21
102:17	<b>closely</b> 139:15	161:5	<b>computing</b> 128:8
<b>checked</b> 86:10	<b>closer</b> 125:12	<b>comment</b> 114:21	<b>concept</b> 11:8
<b>checks</b> 89:15	<b>cluster</b> 166:4	114:23	<b>concierge</b> 57:23
<b>chen</b> 3:17	<b>clusters</b> 67:14,15	<b>commiserate</b>	101:13
<b>chhabria</b> 49:18	68:9 162:14 165:7	27:13,18	<b>conclusion</b> 88:6
<b>choice</b> 93:19	165:10	<b>common</b> 69:2	<b>conclusively</b> 77:10
<b>choices</b> 61:17	<b>code</b> 8:22 83:24	98:11 100:18	<b>concrete</b> 95:16
<b>chose</b> 91:22	84:2,7 86:2 111:9	105:9 110:21	<b>confidence</b> 86:24
<b>chris</b> 3:6 102:22	111:14 114:16,19	133:11 164:9	100:17
<b>circumvent</b> 19:12	114:22,22,24	<b>commonly</b> 150:6	<b>confident</b> 82:11
19:15	115:4,22 116:7,15	<b>communicate</b>	137:25
<b>clarification</b> 4:9	116:23 127:15	172:18	<b>confidential</b> 1:9
169:9,23	133:14,15,18	<b>communications</b>	170:13,13
<b>clarify</b> 49:5	140:12 145:18	90:14	<b>confidently</b> 22:25
102:18 104:23	<b>colin</b> 3:18	<b>companies</b> 76:23	34:24
<b>clarifying</b> 169:11	<b>collect</b> 14:2,14	139:20 171:15	<b>configure</b> 74:23
<b>clarity</b> 20:19	21:3,5,20,25 22:3	<b>company</b> 77:14	<b>confirm</b> 35:23
<b>clark</b> 3:16	22:6,11,19 23:1,3	84:19 86:3,20	<b>confirmation</b>
<b>classify</b> 122:21	23:6 25:24 65:5	104:8	118:3
<b>clear</b> 26:17 32:21	119:24 120:1	<b>compared</b> 83:22	<b>confused</b> 165:9
77:18 79:1 116:3	<b>collected</b> 5:18,21	128:16	<b>connect</b> 22:2
116:21 118:5	24:15,17 25:23	<b>comparison</b> 11:8	23:20 41:18,18
119:19 121:13	30:14 46:5 52:19	<b>compatible</b> 149:10	42:12 44:8 65:9
123:24 141:21	62:9 119:20		128:23 155:2,3

## HIGHLY CONFIDENTIAL

[connected - data]

Page 7

<b>connected</b> 23:18 42:9 45:7 56:24 58:17 <b>connecting</b> 66:3 <b>connection</b> 23:7 105:18,19,22 115:15,18 124:4 <b>connections</b> 60:13 128:18,21 <b>connectivity</b> 23:21 67:20 <b>connects</b> 168:17 <b>consent</b> 108:20,23 108:24,25 <b>considering</b> 124:5 <b>consolidated</b> 137:8,8 <b>constraints</b> 15:7 <b>consumer</b> 1:3 2:3 103:19 133:18,19 139:3,6,10 142:24 <b>contacts</b> 23:16 <b>contain</b> 32:4,5,18 65:25 <b>contained</b> 56:17 99:20 <b>containing</b> 117:9 167:11 <b>contains</b> 118:14 121:9 132:8 <b>content</b> 37:9 64:24 76:6 91:3 103:4 126:16 141:18 142:20 <b>contents</b> 74:21 91:8 <b>context</b> 9:4 93:12 103:1,11 129:14 <b>continue</b> 13:3,5 20:5 35:17 59:11	<b>contract</b> 14:22 16:8 108:13,24 113:17 140:6,12 <b>contracts</b> 111:24 125:22 <b>contractually</b> 109:1 <b>contribute</b> 84:21 <b>control</b> 11:19 23:2 23:3 81:6,10 <b>controlled</b> 9:25 109:1 <b>controls</b> 18:24 <b>conundrum</b> 92:4 <b>convenience</b> 4:20 <b>convenient</b> 163:11 <b>conversation</b> 75:9 80:2 <b>conversations</b> 59:23 121:4 <b>cookie</b> 22:8,14,15 56:23 57:3 58:16 58:16 <b>cookies</b> 22:17,19 56:23 69:21 70:4 70:7 82:21 <b>copy</b> 170:1 <b>core</b> 157:23 <b>corley</b> 49:18,19 <b>correct</b> 82:20 87:7 90:8 108:8 112:2 112:11 136:13 146:22 150:16 <b>correctly</b> 82:19 109:10 156:17 <b>cost</b> 50:2 63:2 64:4 64:4 <b>counsel</b> 7:1 55:17 57:14 125:11 134:11 142:2 155:14 162:4,19	166:18 167:15 173:10 <b>count</b> 129:3 158:7 <b>counter</b> 158:6 <b>country</b> 145:18 <b>counts</b> 6:22 <b>county</b> 173:2 <b>couple</b> 19:24 76:22 166:1 <b>course</b> 8:4 64:12 95:18 <b>court</b> 46:4,5 48:11 49:15,17 51:12,21 64:20 68:1,5 79:1 80:13 92:8 101:22 119:19 123:21 171:19,20 172:3,4 <b>court's</b> 5:7 62:8 125:8 <b>cover</b> 35:2 58:22 100:2 115:10,10 <b>covered</b> 165:7 <b>covering</b> 171:23 <b>covers</b> 81:5 125:7 <b>cow</b> 88:24 <b>cows</b> 88:17,17,17 88:17,22 <b>crazy</b> 167:16 <b>create</b> 44:19 45:13 45:16,17 51:22 58:17 71:6 89:17 93:21 120:15 132:15 <b>created</b> 44:25 45:24 89:8,10 129:18 143:6 149:1 <b>credentials</b> 14:2 14:14,16 15:17 <b>crr</b> 1:24 2:25 173:22	<b>crutcher</b> 3:15 <b>csr</b> 1:24 2:25 173:22 <b>cuff</b> 73:20 166:22 <b>culture</b> 83:22 <b>cumbersome</b> 71:14 <b>curious</b> 106:16 145:8 <b>current</b> 16:1 114:7 132:11 <b>currently</b> 16:5 33:22 <b>custody</b> 72:24,25 73:1,11 <b>custom</b> 76:1,21 91:19 129:11,13 129:19 139:13,19 <b>customer</b> 36:5,11 <b>customers</b> 56:13 <b>customized</b> 139:24 141:5 <b>cut</b> 50:10 159:9 171:2
			<b>d</b>
			<b>d</b> 68:7 133:4 <b>dallas</b> 3:19 <b>data</b> 4:7,13,15,19 4:21,21,23,23,24 4:24,24 5:2,6,13 5:18,21,23 6:3,6 7:6,12,14 9:4,14 10:22,24 11:1,3,5 11:6,12,13 13:18 18:24 19:5,24 20:18,23 21:3,6,20 21:21 22:8,14 23:1 24:7,9,11,15 24:16,21,23 25:1 25:20,22,24 26:14 26:17,21 27:2,4,6



## HIGHLY CONFIDENTIAL

[data - delete]

Page 8

27:17,22 28:2,23	91:2,24 92:13	159:1 161:2	<b>deceptively</b> 46:21
29:8 30:4,14,16	94:5 96:2,17,21,24	162:16 163:9	<b>decide</b> 93:16
31:6,8,10,12 32:5	98:4,7,10,15,16,16	165:12 166:3	<b>decided</b> 142:19
32:14,14,18 34:4	98:19,19,24 99:4,7	167:17 168:10	152:24
34:14,19 35:8,13	99:9,12,15,16,20	<b>database</b> 31:7	<b>decision</b> 36:23
35:15,20,25 36:1,6	102:3,4 103:12,23	38:9,21 43:6 51:7	37:2 64:5
36:9,13,17,25 37:1	105:8,9,12,21	51:8,10 54:10,11	<b>dedicated</b> 151:6
37:14,14 38:1,14	106:2 111:7,13,17	93:18 94:2,25,25	151:13
38:17,23,23 39:3	111:19 112:7	95:20 97:13,23	<b>deep</b> 29:24 76:12
39:13,17 40:2,5,11	114:12 116:1,6,14	98:17,23 110:19	115:22
40:13 41:23,25	117:14,23 118:7	115:15,16 116:20	<b>default</b> 83:14
42:2,8,9 43:7	118:17 119:13,14	122:5 124:16,16	<b>defendants</b> 3:13
44:17 45:1,6,18,23	119:20,21,22	125:7 147:6 153:6	<b>defer</b> 104:3 117:17
46:5,6,7 48:1,4,6,8	120:5,10,11	165:13,15 167:10	138:6
49:15 50:22,25	121:14,14,15,16	<b>databases</b> 38:5	<b>define</b> 51:22 52:1
51:1 52:8,18,19	121:16 122:7,8,14	67:14 68:17 92:11	103:22 125:14,15
53:18,19 56:19,23	122:21,24,24,25	92:12 153:7	<b>defined</b> 5:13,16
56:24 57:2,3	123:1,2,5,10,11,13	158:20,21 171:24	6:7 49:15 64:21
58:16 59:15,16,19	123:22,23 124:1	<b>dataset</b> 49:4 97:7	74:22 125:8 139:4
59:22 60:3,5,7,13	124:12,19,21,23	<b>datasets</b> 56:14,17	140:16
60:17,22 61:24	124:24 125:3,17	118:20	<b>defines</b> 111:6
62:9,11,13,17 63:7	125:23 126:22,24	<b>date</b> 15:25 33:21	116:6,8,11 123:21
64:13,15,19,23,24	127:22 128:18,21	38:19,21 40:20	<b>defining</b> 124:25
65:2,4,6,17,21,21	128:23 129:21,24	41:2 76:11 153:4	<b>definitely</b> 73:23
66:4,8,11 67:9	130:2,7,13,17,22	<b>dated</b> 173:15	78:4 99:3 104:5
68:12 69:4,10,12	131:24,25 132:15	<b>daughter</b> 167:24	117:15
69:14,19 70:17,20	133:11,24,25	<b>davis</b> 3:10,18	<b>definition</b> 5:1,7,12
70:22,23,25 71:4,6	134:5,18,21,23	<b>day</b> 38:1,1 41:4,15	5:16,25 32:21
71:7,15,16,22	136:4,8 137:25	50:19,19	46:3 49:14 51:12
72:12,16,21,23,25	138:1 141:24	<b>days</b> 54:8 68:17	51:18 62:8 99:11
73:1,7,12,12,13,24	143:1,5,5,15,21,25	113:19 170:20	103:5 116:13
74:3,7 75:4,14,16	144:2,8,10,11,13	<b>db</b> 30:3 67:12	120:18 122:14
77:8 78:16,19	144:18 145:5	153:22 168:23	150:10
79:4,23 81:16,17	146:20 147:7,12	<b>de</b> 4:24 103:12	<b>definitions</b> 6:15
81:20,21 82:8,9,15	149:13,25 150:4	<b>deal</b> 121:13	7:15
83:6,9,16,18,20	150:12,12,14,19	150:14	<b>definitive</b> 117:21
84:9,24 85:8,21	150:25 151:14,17	<b>dealing</b> 97:9	<b>degree</b> 79:10 96:8
86:20 87:2,2,3,16	151:23 152:7,10	<b>debugging</b> 129:25	98:10 113:7 149:9
87:17 88:9 89:1,9	152:12,16,18	<b>decade</b> 167:18	<b>delete</b> 103:23
89:13,17,18,18	153:24,25 157:22	<b>decent</b> 70:17	105:18
90:13,15,17,20,23	157:25 158:4,23		

## HIGHLY CONFIDENTIAL

[deleted - documentation]

Page 9

<b>deleted</b> 103:21	<b>designed</b> 31:4	112:13	<b>directional</b> 126:24
<b>deletes</b> 105:17	40:13 41:1,4	<b>die</b> 32:9	<b>directly</b> 39:3,14
<b>deleting</b> 105:11	87:13	<b>difference</b> 82:18	40:3,5 50:16
<b>deletion</b> 102:3	<b>detailed</b> 171:15	129:10 154:20	64:23 67:8 76:19
103:12 104:4	<b>details</b> 69:23	<b>differences</b> 96:9	98:10 111:12
121:14,20	70:14 75:19	115:12	151:23 163:11
<b>delimiter</b> 123:25	<b>determination</b>	<b>different</b> 9:16	<b>disappear</b> 163:10
124:1	49:18	13:24 15:15,16	<b>disclaimers</b>
<b>delineates</b> 145:1	<b>determine</b> 4:10,16	18:14,16 21:17	108:15
<b>delineation</b> 114:25	5:4 7:10 20:9	28:11,12,22 34:21	<b>disclosed</b> 7:11
<b>delivered</b> 41:9,16	90:22 134:20	36:13 38:21 40:11	<b>disconnected</b> 59:4
42:12 65:8	<b>determined</b> 79:3	47:24 50:25 56:16	<b>discoverable</b> 5:13
<b>delivery</b> 41:14	<b>determines</b> 111:13	62:7 63:19 80:7	<b>discuss</b> 4:7 105:5
<b>depend</b> 83:15 85:9	<b>develop</b> 83:19	84:20 90:14 93:22	<b>discussed</b> 6:25
91:16,21 130:24	171:3	99:23 109:9 112:9	56:3 153:17 157:4
152:20	<b>developed</b> 139:12	113:8 123:2	166:13 171:12
<b>depending</b> 10:13	<b>developer</b> 14:9	125:14 128:11	<b>discussion</b> 97:20
21:9 44:24 151:7	18:17 19:8,12,14	131:7 136:16	<b>disneyland</b> 155:11
152:23	21:23 37:12 74:22	137:22 139:4	<b>display</b> 108:12
<b>depends</b> 66:25	79:14 80:17 89:12	140:16 142:11	<b>displaying</b> 108:17
116:15 125:20	89:25 90:2 120:11	143:15 144:13,13	<b>distinction</b> 9:21,22
<b>depo</b> 142:16	123:18,20	149:4,10 152:11	10:2,13 28:1
<b>depos</b> 90:10	<b>developer's</b> 11:19	152:11,22,23,23	122:25 124:10
<b>deprecated</b> 33:22	109:16	154:14 156:24	141:14 157:1
<b>deprecation</b> 70:24	<b>developing</b> 60:7	162:14,15,16	<b>distinguish</b> 56:10
<b>depth</b> 79:8	<b>development</b>	<b>differently</b> 73:6	<b>distributed</b> 118:19
<b>derek</b> 3:5 103:10	133:20	139:5 144:8	<b>divorced</b> 78:17
<b>derive</b> 91:4 118:18	<b>device</b> 9:25 10:21	<b>difficult</b> 43:19	<b>diy</b> 29:1 30:10,11
<b>derived</b> 99:17	10:22 15:7,9	65:17 66:5 77:20	31:14,19,24 34:5
119:5 133:11	23:20 109:17	92:3 99:15 115:21	46:1,2,9 47:5 48:7
<b>derives</b> 118:21	143:15 145:17	127:12 144:19	49:2 50:25 51:9
<b>describe</b> 103:11	<b>devices</b> 21:21 80:8	<b>difficulties</b> 150:24	88:2 90:16,17
<b>described</b> 11:25	144:14 145:7	<b>dig</b> 116:23 155:7	117:13,19 118:8
12:2 16:17 17:3,7	<b>diagram</b> 71:4 73:7	<b>digging</b> 85:13	118:22 119:14
17:17 18:6 96:13	73:12 83:19,20	<b>direct</b> 90:5 96:11	120:4 121:7 132:1
105:7	84:10	<b>directed</b> 105:2	132:3,4 135:3
<b>description</b> 116:21	<b>diagrams</b> 71:7,15	<b>directing</b> 104:25	<b>dloeser</b> 3:8
<b>descriptors</b> 115:4	87:2,2,11,15	<b>direction</b> 54:15,16	<b>document</b> 83:25
147:20	103:25	59:23 60:12,21	<b>documentation</b>
<b>design</b> 41:1 83:24	<b>dialogue</b> 108:5,11	145:2 173:9	104:1 106:7
87:11,12	108:18 111:3		114:10,11,12,16



## HIGHLY CONFIDENTIAL

[documentation - engage]

Page 10

115:3 117:2,3,5,7 117:10 118:8,13 119:10,15 153:1,3 153:4 <b>documents</b> 24:3,9 102:21 103:11 <b>doing</b> 88:4 89:17 107:6 149:11 <b>doll</b> 13:7 <b>dollars</b> 30:22 <b>domain</b> 22:18,20 22:23 <b>domains</b> 22:20,22 23:1,3 <b>dot</b> 82:18 112:13 167:10,10 <b>download</b> 56:15 <b>downstream</b> 50:8 <b>dozen</b> 92:22 <b>draw</b> 124:14 <b>drawing</b> 56:14,16 122:25 168:16 <b>drop</b> 39:23 101:3 <b>dropbox</b> 139:21 140:2,3 <b>dropped</b> 61:24 68:13,15 72:3 <b>ds</b> 71:18 <b>due</b> 15:6 61:16 69:13 <b>dumbo</b> 42:21,22 42:23 43:1 155:9 167:23 <b>dump</b> 97:6 <b>dumps</b> 124:22,23 <b>dunn</b> 3:15 <b>duplicate</b> 98:9 99:4 135:16 138:3 <b>duplicated</b> 45:2 98:3	<b>duplicates</b> 46:14 <b>duplication</b> 44:17 99:23 <b>duplicative</b> 30:5,8 30:10 51:3 142:3 142:4,5,22 143:3 143:17 151:19 <b>dyi</b> 117:23 118:1,9 118:13,16,19 119:6 120:15,17 120:19,23 132:13 138:5,23 140:13 141:22,23 142:22 143:17,20 144:17 145:16 151:18,19 159:25	129:17 140:11 145:11 148:21 155:13 161:22 166:10,17 <b>elastic</b> 33:25 34:2 <b>elasticsearch</b> 155:18 <b>elements</b> 116:2,7 145:5,18 <b>elia</b> 7:20 8:14 16:4 16:9 17:11,24 23:14,23 32:1 34:23 35:3,5,16 43:2,18 44:4 48:16 49:7 50:12 55:2 57:17,20,21 59:11 64:9,18 65:3,14,24 66:3,10 66:17 69:9 70:5,6 70:10 72:4,5,10,14 72:18 74:18,19 75:12 76:24,25 79:9,10 80:4,17,21 83:1,3 85:16 86:1 87:6,7,18 92:20,21 93:1,9,11,19 94:4 94:9,19 95:8,15 96:5,8,19 97:2,24 98:8 100:7,13,16 101:3 102:14 104:25 105:1,3 106:1,12,19,22 107:11 109:10,21 109:25 110:8,14 110:21 111:8 112:3 113:4,13,19 113:24 114:5,14 114:23 122:12,13 122:19 123:3,8,15 124:3,7 125:19 126:5,13,23 127:2	127:6,23 128:10 128:13 129:8,12 131:2,3 133:6 135:24 136:1,17 137:4,13,24 140:21,22 141:2,3 141:10 143:24 144:2,6 145:12,14 145:22,23 148:21 148:23 149:3,8,15 149:19,22 150:1,6 150:13,18,22 153:11 154:7,9,12 154:19,21 155:12 155:22 156:10 157:9,12,15,18 158:18 159:18 160:10,25 161:6 161:11,20,21 162:1,24 163:5,16 163:21 164:4,21 165:3,17 166:12 166:17 168:6,13 170:24 <b>embedded</b> 15:8 59:16 64:14 108:11 <b>eminent</b> 95:13 <b>enable</b> 40:5 139:13 <b>enabled</b> 140:6 <b>enables</b> 40:2 44:18 <b>encoded</b> 148:5 <b>encryption</b> 80:22 <b>ended</b> 27:3 63:4 <b>ends</b> 29:1 41:21 49:1 70:22 77:17 83:14 85:8,10 133:3 <b>engage</b> 28:25
	<b>e</b>		
	<b>e</b> 1:23 2:25 39:25 68:6 76:14,18 173:4,22 <b>earlier</b> 54:13 131:10 132:10 150:23 <b>early</b> 70:23 170:21 <b>earth</b> 30:17 <b>easier</b> 60:25 89:8 121:25 <b>easy</b> 114:6 <b>edition</b> 40:21 <b>effective</b> 92:8 <b>effectively</b> 83:24 111:20 <b>effects</b> 112:15,16 <b>efficient</b> 92:8 <b>effort</b> 47:1 63:17 104:17,18 118:24 <b>either</b> 8:20 27:20 32:1 35:6 42:24 44:10 47:6 48:25 70:3,11 76:4 82:7 100:4 110:11		

## HIGHLY CONFIDENTIAL

[engagement - example]

Page 11

<b>engagement</b> 49:5 143:22	<b>essentially</b> 113:16	87:4,10 88:14	162:23 163:2,3,7
<b>engine</b> 126:7	<b>established</b> 134:19	89:3,14 90:8 91:6	163:15,19,25
<b>engineer</b> 84:8 86:3	143:18 152:1	91:16 92:1,3,16	164:7,15,24 165:5
87:1 88:21 90:4	<b>et</b> 21:1 125:16	94:4,12 95:2 97:1	165:14,17,24
90:13 97:12 127:5	<b>etsy</b> 125:16	97:8 98:21 99:2	166:11,16 167:8
127:8 138:4,9,13	<b>eu</b> 121:13	99:22 100:6,12,15	168:3 169:1
138:21,23 151:18	<b>eugene</b> 9:20 10:9	102:11 104:3	<b>european</b> 103:18
<b>engineer's</b> 119:17	10:17,19 11:16	106:11,13 107:13	<b>evening</b> 16:3
<b>engineering</b> 83:22	12:1,6,12,17 13:8	107:25 108:6,22	<b>events</b> 74:23
83:23 85:6 98:24	14:6,20,25 15:12	109:5 110:10	<b>eventual</b> 29:17
104:16,19 105:5	15:19,24 17:9,19	111:18,22 112:4	<b>everstore</b> 93:2
133:21 141:8	18:7,10,18,23 20:2	112:22 113:3	94:8,9,14,23
151:18	21:3,7,9,13 22:5,9	114:15 115:2,9,20	155:24 156:14
<b>engineers</b> 5:9 8:23	22:17,25 23:11	116:3,8 117:3,17	168:24
42:15,18 49:8,24	24:19 25:9 26:24	118:11,23 120:13	<b>everybody</b> 5:17
60:1 77:24 85:17	27:13,18,25 28:13	121:8,19,25 122:9	58:10 145:25
87:9 90:4 138:9	28:17 29:11 31:1	127:25 128:20,24	148:17 170:1
138:20 143:16	31:18 32:19,20	129:3,6,23 130:1,8	<b>evolved</b> 15:22
146:17 151:18	33:1,10,14,21 34:6	130:19,24 131:4	<b>evolving</b> 153:5
152:24	36:8,15 37:7,16,25	131:18 132:2,6,10	<b>exact</b> 8:25 81:21
<b>entertain</b> 102:8	38:8,13,19 39:10	132:17,20 133:1,5	108:23 115:10
<b>entire</b> 25:1 34:1	39:16 40:9,18	133:10,15,22	136:15 158:24
63:21 70:23 88:19	41:11,20 42:14,24	134:3 135:4,20	<b>exactly</b> 19:7 21:13
111:9 138:7	43:19 44:2,9,13,16	136:7,11 137:12	23:2 29:13 38:8
<b>entirely</b> 91:21	45:16,19,25 46:10	137:18 138:6,12	77:7 85:9 108:14
113:13	46:24 47:8,11,17	138:22 139:7,22	108:22 109:5
<b>entirety</b> 40:23	47:23 51:18 52:4	140:10,19 141:1	115:22 116:15
113:5 127:15	53:3,8,17 54:4	141:11 144:1,16	117:10,18 120:14
<b>entity</b> 52:3	59:18 60:3,9 61:5	145:13 146:23	151:21 157:24
<b>ephemeral</b> 161:2	61:12,15,22 62:2	147:14 148:3,10	158:3 167:23
<b>epically</b> 34:12	62:19 65:14 66:21	148:24 151:3,21	<b>example</b> 10:17,19
<b>equality</b> 92:6	66:25 67:6,13	152:3,6,13,19	11:4 23:16 24:1
<b>error</b> 59:4 128:15	68:11 69:1,12,22	153:2,9,18,23	29:7,15 31:6 32:5
<b>especially</b> 136:5	70:13,19 71:5,21	154:5,11,15,23	32:23 34:15 36:19
<b>espinoza</b> 57:23	72:18 73:3,15	155:4,10,19,25	36:19 41:24 42:10
101:2,13	74:8,13 75:10,17	156:5,8,12,17,23	42:13 43:8,17,21
<b>esq</b> 3:5,5,6,6,9,10	75:25 76:9 77:6	157:7,11,14,17,20	43:22 44:10 48:23
3:10,16,16,17,17	78:7,21 79:7 80:7	157:24 158:6,11	51:15 53:11 57:1
3:18,18	81:4,10,19 82:11	158:17,20,24	60:25 63:6,16
	82:15 83:9,13,21	160:9,12,23	65:3,8 68:16
	84:3,14 85:5 86:6	161:16,25 162:9	71:19 81:23,24

## HIGHLY CONFIDENTIAL

[example - familiar]

Page 12

82:6 86:1 89:20 90:12 95:16 99:5 99:5 110:17 116:9 116:10,17 120:5 122:20 126:5 131:8,12 133:23 134:4,24 135:20 140:1 141:4 144:3 144:10 150:8 161:3 164:9,16 <b>examples</b> 135:25 <b>excellent</b> 166:8 <b>exception</b> 87:4 148:4 <b>excessive</b> 7:13 <b>exchange</b> 15:1 108:16 <b>excited</b> 167:25 <b>execute</b> 111:14 <b>exemplar</b> 56:25 58:18 <b>exempt</b> 84:24 <b>exercise</b> 120:23 127:7 <b>exhibit</b> 141:16 142:16,16 144:25 <b>exist</b> 31:17 62:21 66:12 87:14 100:1 132:14,21,22 136:21 137:3 141:15,23,23 <b>existed</b> 27:23 96:17 127:10 <b>existence</b> 88:19 <b>existing</b> 99:12 172:15 <b>exists</b> 46:25 63:25 73:17 129:21 141:22 143:5 <b>expect</b> 24:22 35:20 65:25	<b>expected</b> 35:12 <b>expensive</b> 95:21 96:1 <b>experience</b> 26:8 60:14 95:12,14,25 125:21 134:13 139:25 161:18 <b>experiences</b> 139:4 139:6,10,14,16,20 142:24 143:2 <b>expert</b> 29:24 <b>expertise</b> 8:9,12 104:5 140:20 <b>explain</b> 36:1 71:4 71:16 73:7 103:1 <b>explaining</b> 73:13 <b>explicit</b> 49:17 113:10 114:18,24 142:17 <b>explorer</b> 17:24 18:1,3,5,8 148:22 149:15 <b>expose</b> 84:19 <b>externally</b> 56:9 <b>extract</b> 89:12 90:13,17 147:25 <b>extracted</b> 47:6 48:25 85:24 88:10 89:1 95:24 <b>extracting</b> 49:3 90:20 <b>extremely</b> 170:25	20:24 21:25 22:20 22:21,23,24 23:6,8 23:18 25:23 28:2 31:22 32:3,16,17 34:3 36:4 38:1 39:3,9,13,14 40:6 43:10,10 45:13 46:3,19 47:4 48:10,24 50:23,24 51:16,17 52:2,23 52:24,25 53:2,4,22 53:24,25 54:1 55:17,18 56:8,11 56:12,17 57:3,9,10 58:18 59:15 60:14 61:7 64:13,16,16 69:9,10,14,16 70:9 70:16 71:22 72:16 73:1,8,11 74:6,22 75:4,6,8,14,16 77:2 79:4,11,13 82:8,18 88:21 90:2,3,4 95:19 100:24 101:10 102:2,2,9,17 103:22 106:6 107:15 108:3,20 111:23 112:24 113:1 118:20 123:14 126:10,16 126:18,21,24 127:20 128:17,18 129:18 134:17,25 135:2,13 139:12 139:15 140:4 141:19 142:2,20 143:2,22 144:23 145:1 146:13 151:22,24 155:15 155:15 159:4,7 160:15 166:19	167:3,10 169:17 171:19,24 <b>facebook's</b> 11:22 12:22 13:4 14:11 56:13 80:24 128:23 143:22 <b>facing</b> 57:9 62:20 133:19 <b>fact</b> 29:16,19 63:25 85:17 111:22 142:14 <b>failed</b> 34:12 <b>failure</b> 97:10 <b>fair</b> 31:22 36:12 42:15 90:11 124:9 130:1 <b>fairly</b> 98:5 171:3 <b>falconer</b> 3:16 7:1 7:19 8:3,5,7 24:2 34:22 50:9 54:20 54:23 55:2 57:14 57:16,22 58:3 73:18 89:23 90:7 101:13 104:6,13 117:20 119:3 125:11 133:4 134:11,14 138:11 138:14 142:6 143:8 159:14,24 160:19 162:18 166:21 169:19 <b>fall</b> 159:10 <b>falls</b> 153:4 159:11 <b>familiar</b> 9:20 33:3 48:19 50:18 52:14 70:7,14 75:19 96:9 100:13 107:11 110:22 113:24 125:25 126:3 127:23 142:12 144:7
	<b>f</b>		
	<b>face</b> 19:17 81:18 133:18 <b>facebook</b> 1:3 2:3 4:8,12,12,18 5:9 6:11 7:1,13,18 10:20 11:5,6 14:2 14:3,17 15:2,8 17:8 19:3,4 20:18		

## HIGHLY CONFIDENTIAL

[familiar - ftp]

Page 13

148:22 155:12,22 157:9 160:10 161:21 <b>familiarity</b> 93:6 96:12 <b>fan</b> 60:15 <b>far</b> 56:3 63:8 115:2 120:3 121:4 121:25 145:16 162:18 <b>fashion</b> 49:1 64:20 65:19 66:13,16,20 105:11 <b>fast</b> 171:5 <b>faster</b> 38:22 <b>feature</b> 17:6 <b>february</b> 1:14 2:24 4:2 7:22 173:15 <b>feed</b> 117:7,8 131:12 <b>feel</b> 5:10 8:21 21:17 55:19 155:16 171:1 <b>fellow</b> 172:13 <b>fetch</b> 11:6,12 <b>fewer</b> 163:6 <b>field</b> 111:10 <b>fields</b> 93:22 114:18,25 130:12 147:19,25 148:4 <b>figure</b> 4:17 25:14 30:2,3,13 45:9 48:10 50:21 74:2 77:3 84:2,6,9 92:7 98:20,22 119:11 120:9 136:3 139:1 143:19 <b>figuring</b> 51:4 114:10	<b>file</b> 30:10,11 31:19 31:25 34:5 46:9 47:5 48:7 49:2 50:25 51:9 88:2 90:16,18 114:22 118:8,9,13,16,19 118:22 119:6,14 120:4,15,17,23 132:1,3,7 135:3 141:22,23 143:17 143:20 144:18 145:16 151:20 <b>filed</b> 7:23 91:1 <b>files</b> 94:10 120:19 127:16 132:13 142:22 145:4 <b>final</b> 89:6 <b>finally</b> 5:10 <b>find</b> 13:17 23:17 40:21,25 50:15 56:7 57:20 60:17 61:1,2 63:17,20,24 70:20 71:7 77:21 77:23,25 78:23 81:16 82:2,3 84:22 85:6 87:20 89:20 91:23 94:15 96:24 99:15 115:22 120:17,20 120:24,24 127:5 141:12 143:10,12 148:11 165:22 166:22,23 167:19 <b>finding</b> 86:8 <b>fine</b> 16:10 17:13 24:7,8 35:1,2 104:12 147:18 <b>fingertips</b> 138:16 <b>finish</b> 16:24 125:12 148:17	<b>firewall</b> 127:22 <b>first</b> 5:18 9:12 11:17,24 30:5 46:16 56:7 59:8 84:4 88:12 103:16 120:19 146:19 <b>fit</b> 120:12 <b>fits</b> 158:4 <b>five</b> 55:7,10,14,15 55:20,21 100:23 101:1,5 145:21,25 159:6 170:20 <b>flag</b> 24:12 <b>flip</b> 15:10 <b>flow</b> 8:16 57:2 70:23 71:7,10 80:12,16,18 87:2 152:23 <b>flows</b> 71:4,16 73:13 77:8 86:20 87:16 <b>focus</b> 24:16 26:20 92:25 171:4 <b>focused</b> 148:19 <b>folded</b> 136:14 <b>folder</b> 140:3 <b>folks</b> 159:19 <b>follow</b> 17:15 20:15 31:23 55:18 56:2 56:5 88:13 104:22 115:7 121:17 141:11 147:15 172:3 <b>followed</b> 137:20 <b>following</b> 89:24 106:17 153:13 <b>fonti</b> 3:9 <b>forget</b> 16:25 24:16 <b>forgot</b> 91:11 <b>form</b> 35:20 76:15 89:4 102:5 118:1	<b>formal</b> 118:3 170:8 <b>formally</b> 119:4 <b>format</b> 68:12 <b>forms</b> 93:22 156:6 <b>forth</b> 36:3 173:7 <b>forward</b> 9:1 59:11 103:17 <b>four</b> 26:9 149:4 168:22 <b>fql</b> 109:19,22 137:15 138:2 <b>frame</b> 16:14 <b>framing</b> 56:5,7,21 <b>frankly</b> 127:15 <b>free</b> 55:19 87:22 89:4 155:16 <b>freely</b> 105:16 <b>frequently</b> 50:16 <b>friend</b> 12:18,19 93:15 109:19,23 110:1,3,5 164:11 164:17 <b>friends</b> 11:4,23 12:3,8,9,13,14,18 13:6,6,6,9,14 23:17 106:18,19 106:20,22 107:5 109:8,8,12,13 126:11,11 168:9 <b>friendship</b> 107:3 <b>friendships</b> 106:24 <b>front</b> 87:1 116:22 134:6,9 156:6 160:2 <b>fronting</b> 45:3 99:12 <b>froze</b> 80:14 <b>fruitful</b> 104:14 <b>ftp</b> 129:5
--	--	--	---

## HIGHLY CONFIDENTIAL

[fulfilled - getting]

Page 14

<b>fulfilled</b> 131:15	40:17 41:6,14	109:18,22 110:6	158:14,19,22
<b>full</b> 12:21 15:7	42:4,18,25 43:4,24	110:12,18 111:2	159:3,23 160:5,11
93:11 103:14	44:5,12,14 45:4,17	111:16,21 112:19	160:14,20 161:4,8
149:22	45:20 46:2,15	113:1,11,16,23	161:13,19,23
<b>fully</b> 8:23 10:7,8	47:3,9,15,21 48:2	114:2,9,21 115:5	162:2,17,21
171:17,18	48:20 49:12 50:11	115:14,24 116:5	163:12,17,22
<b>function</b> 83:12	50:20 51:20 52:10	117:1,12 118:4,15	164:6,12,19,22
111:11 116:1	53:7,14,20 54:17	119:1,7 120:13,25	165:1,4,8,22 166:7
<b>functionality</b> 26:7	54:20,22,25 55:5	121:12,23 122:3	166:13,18,25
74:20 76:3,12,13	55:25 57:11,18	122:12,17,23	167:12,21,24
76:17 96:15	58:1,7,11,13,19	123:6,12,19 124:6	168:5,12,15 169:5
<b>funny</b> 61:12	59:6 60:1,4 61:3,9	124:9 126:2,8,21	169:17 170:3,6
<b>further</b> 4:9 69:7	61:14,19 62:1,5	126:25 127:4,21	171:7,17,22 172:9
169:9,13 173:10	64:7,11 65:1,5,23	128:4,12,17,22	172:17
<b>future</b> 24:22	66:2,7,14,23 67:3	129:1,4,10,16,25	<b>gate</b> 111:9
<b>g</b>	67:10,20 68:2,7,22	130:4,14,20 131:1	<b>gated</b> 113:5,13
<b>gae</b> 17:18,20	69:5,17 70:1,8,12	131:16,22 132:4,9	127:11
149:12,13,24	70:15 71:3,9 72:4	132:13,18,24	<b>gdpr</b> 102:6 103:15
150:10	72:8,13,15,23 73:5	133:2,7,13,17	103:18 121:15
<b>gain</b> 110:15	73:22 74:10,17	134:2,10,16	<b>gender</b> 65:6
<b>gaps</b> 120:17	75:1,13,20 76:7,22	135:22 136:2,9	<b>general</b> 6:14 27:5
<b>garrie</b> 4:4,6 8:2,4	77:1 78:3,10,25	137:2,7 138:4,8,20	28:18 35:7 56:4,7
8:6,17 9:10 10:3	79:9,21 80:19	138:25 139:9	59:14 62:2 136:23
10:16,18 11:14,21	81:2,9,15 82:5,14	140:8,11,21,24	137:13
12:5,11,15,20	82:25 83:5,11,18	141:2,7,13 142:13	<b>generalized</b> 62:21
13:20 14:8,24	84:1,5 85:1,14	143:14 144:4,12	63:2
15:10,13,20 16:2,7	86:16 87:6,8 88:1	144:21 145:15,24	<b>generally</b> 37:16
16:10,12,19,22	88:16 89:11 90:1	146:4,17 147:10	44:21 52:6,7
17:10,13,21,25	90:3,9 91:10,23	147:16,22 148:7	59:20 60:9 65:17
18:9,13,22 19:6	92:2,5,18,24 93:7	148:13,16,25	67:8 87:17 110:15
20:4,14 21:5,8,11	93:10,13 94:1,7,16	149:7,12,17,20,24	113:25 133:22
21:16 22:7,12,21	94:24 95:4,11,23	150:4,9,16,19	142:2 144:9 153:4
23:4,19 24:2,6	96:6,16,23 97:5,15	151:1,16,22 152:4	<b>generate</b> 83:23
25:8,11 27:8,15,21	98:4,14 99:1,18	152:8,15,25 153:6	115:3,6 118:7
28:3,16,24 30:1	100:4,8,14,20	153:12,20 154:2,7	120:19,23
31:13,21 32:2,21	101:9,16,19,23	154:10,13,18,22	<b>generic</b> 57:8 77:13
32:24 33:5,13,18	102:6,16 103:14	154:25 155:9,11	<b>getting</b> 16:20 29:5
33:24 34:8 35:1	104:12,15 105:1	155:14,24 156:2,7	35:19 37:24 70:25
35:14,21 36:10	105:23 106:3,14	156:11,15,20	71:1 74:2 79:7
37:4,13,17 38:6,10	106:20 107:9,23	157:3,8,10,13,16	84:10 89:16
38:16 39:2,12,22	108:4,19 109:3,6	157:19,22 158:3,9	107:25 142:11



## HIGHLY CONFIDENTIAL

[getting - header]

Page 15

166:5 <b>gibson</b> 3:15 <b>gibsondunn.com</b> 3:20 <b>gist</b> 14:5 58:20 <b>give</b> 20:15 32:5 59:2 65:7 70:16 71:18 85:11 90:12 95:15 104:11 105:3 134:24 166:9 <b>given</b> 7:14 11:11 43:19 67:16 94:13 105:20 111:14 114:17 129:14 131:5 146:15 148:10 155:21 171:11 <b>gives</b> 25:25 59:4 <b>glad</b> 5:9 10:8 <b>gladly</b> 155:17 <b>go</b> 4:5 8:6 19:4,8 25:13 26:21 27:10 28:7 32:3 34:25 40:17,24 48:10 50:21 53:19 54:14 55:10,11,12,14,23 57:2,4,14 58:11,23 59:6,10 61:1 63:19,21,24 66:2 78:6 81:9 84:12 85:4 91:2 92:17 93:1 94:20 97:7 101:5,7,21 113:3 116:18,19 121:3 123:7 127:14 133:9 134:25 136:24 137:12 138:11 143:21 146:1,5 169:22 170:3,6 172:19	<b>goal</b> 4:16 5:4 7:5 <b>goes</b> 28:19 68:24 75:21 83:8 85:22 <b>going</b> 7:16 16:14 20:16,22 35:13 40:23 50:1 51:15 54:21 55:23 56:15 59:13 64:2,8 78:17 80:11,15,18 83:12 85:22 87:3 92:10 101:20 102:2 106:4 108:14 125:9 133:10 138:12 142:25 146:5 148:4 151:3,5,6 156:25 165:24 166:4 167:25 170:10 171:5 <b>good</b> 16:3 25:21 27:21 31:22 44:11 44:15 46:10 64:6 65:13 89:17 99:5 100:8 101:21 102:17 104:19 119:7,8 121:1,2,20 126:8 131:23 145:16 <b>google</b> 34:9,11,11 34:13,14,14,17,17 <b>gotta</b> 20:21 67:23 <b>gotten</b> 162:19 <b>gould</b> 3:6 <b>gps</b> 22:2 <b>grab</b> 38:8 <b>grant</b> 37:2 107:10 108:11,16 112:12 113:21 <b>granted</b> 12:10,18 18:17 97:25 111:23	<b>granting</b> 23:22 113:21 <b>grants</b> 18:18 114:7 <b>graph</b> 11:22 12:23 12:25 13:1,5 17:21,24 18:1,3,5 18:8 25:1,15 52:13 74:9 82:2 82:24 121:7,8 122:8,14,16,22 123:4 135:6,7 136:6,9,12,18 137:8,10 138:2 144:15 148:20,22 149:1,14,15,16,21 149:25 150:1,2,2,3 150:10,11,13,14 150:20 168:17 <b>graphical</b> 17:23 <b>great</b> 25:16 48:22 63:16 142:3 <b>green</b> 88:17 <b>ground</b> 39:24 <b>group</b> 6:6 36:11 48:21 49:24 124:18 <b>guarantees</b> 163:6 <b>guard</b> 114:3 <b>guess</b> 9:5,12 17:8 19:16 39:19,25 45:11 52:11,22,25 65:7 68:10,22 74:5 75:3 92:17 102:25 103:16 106:10 109:7 117:12 121:5 129:4 134:11 140:13 144:19 151:17 152:15 162:9 167:8,9 168:10	<b>guesses</b> 92:20 <b>guts</b> 142:7 <b>guys</b> 5:10 17:22 18:2 22:3 37:21 39:5 52:15 70:3 71:3 91:1 112:21 145:20 148:21
			<b>h</b>
			<b>habit</b> 125:5 <b>half</b> 92:22 <b>hand</b> 132:11 155:16 162:13 <b>handful</b> 124:17 <b>handle</b> 31:9,11 100:2 122:11 134:7 <b>handled</b> 38:14 81:6 <b>handles</b> 73:8 84:18 <b>happen</b> 70:23 73:16 112:15 <b>happened</b> 158:8 <b>happens</b> 28:18 85:23 102:18 115:11 116:14 <b>happy</b> 55:9 104:9 119:3 138:16 159:15 160:1,4 166:22 <b>hard</b> 46:13 49:22 77:1 165:14 <b>harder</b> 113:9 <b>hash</b> 76:17 <b>hashing</b> 76:7 <b>hazy</b> 108:1 <b>head</b> 73:19 107:12 110:9 127:3 149:4 <b>header</b> 71:24 72:2 73:24

## HIGHLY CONFIDENTIAL

[hear - includes]

Page 16

<b>hear</b> 20:9 48:14 55:9 67:25 68:1 <b>heard</b> 19:13 31:22 44:2 80:14 102:14 <b>hearing</b> 1:13 2:23 4:7 8:1 36:2 135:23 170:12,12 170:15,25 172:7 172:10 <b>hearings</b> 171:12 <b>held</b> 1:11 <b>hello</b> 16:2 <b>help</b> 10:17 14:6 32:20 67:16 70:20 166:6 <b>helpful</b> 8:8 16:13 24:2,5,13,20 33:20 36:1 134:11 156:21 160:15 166:8 169:7,23 171:1 <b>helps</b> 8:15,16 53:1 120:20 <b>hereunto</b> 173:13 <b>high</b> 85:2 123:9 130:17 134:1 160:25 <b>higher</b> 99:25 128:14 <b>highlight</b> 43:25 <b>highly</b> 1:9 63:7 64:3 137:25 170:13 <b>history</b> 98:1 127:15 <b>hits</b> 39:24 <b>hive</b> 27:2,5,12 28:18,19 29:17 30:15 31:12,17,19 47:6 48:25 83:10 83:14 84:17 85:22	85:24 88:10,18 89:2,3,13,18 90:5 90:14 91:9,14,14 91:15 92:16,18 99:12 110:22 129:20 130:2,6,10 132:15 133:11,12 133:14,16,24,25 134:7 146:13,20 147:2,11,19 168:20 <b>hm</b> 134:14 <b>hobbies</b> 139:19 <b>hold</b> 36:2 39:23 148:15 <b>home</b> 44:6,7 166:14,15,19 167:5 <b>honest</b> 58:2 <b>honestly</b> 60:3 163:3 <b>hope</b> 16:3 <b>hopefully</b> 59:11 101:2 <b>horrible</b> 34:10 <b>hour</b> 54:23 <b>house</b> 22:4,10,13 92:15 <b>housed</b> 92:11,13 <b>hp</b> 72:2 <b>ht</b> 83:9 <b>http</b> 22:18 71:24 83:10,14 130:2 <b>huge</b> 30:21,23 31:2 <b>human</b> 52:2,3 118:1 <b>hundreds</b> 97:9 <b>hurricane</b> 161:5 <b>hypothetically</b> 34:18	<b>i</b> <b>i.e.</b> 63:7 <b>ian</b> 3:17 <b>ibu</b> 161:23 162:8 <b>ibubiz</b> 161:24 162:8 <b>idea</b> 9:1 26:1,2,3 30:20 42:7 48:15 70:17 88:20 95:24 98:7 100:5 162:4 166:5,8,19 <b>identifiable</b> 141:18 142:18,19 158:13 <b>identification</b> 143:4 <b>identified</b> 4:12 25:13,15 35:24,24 36:7 43:12,22 136:19 143:2,3 145:19 159:19 162:5 165:11 <b>identifier</b> 41:13,17 54:11 57:5 58:15 60:11 61:6 94:13 146:14 <b>identifiers</b> 41:21 41:25 56:6,22 63:21 76:2,4,10 94:22 103:3,8 <b>identify</b> 33:18 66:14 102:23 103:4 116:1 131:23 134:11 147:11,20,21,24 160:1 170:21 <b>identifying</b> 86:7 143:19 150:24 <b>ids</b> 22:15 65:25 76:13 94:13,20 98:13 102:5	<b>images</b> 156:1 <b>imagine</b> 23:25 78:16 141:5 <b>immediate</b> 133:25 <b>implement</b> 38:6 <b>implementation</b> 64:4 116:22 152:20 <b>implemented</b> 78:12 116:16 117:11 126:6 165:19 <b>implementing</b> 77:24 152:24 <b>implication</b> 117:6 <b>import</b> 23:16 <b>important</b> 6:18 142:25 <b>impossible</b> 15:7 42:3 66:5 <b>impressed</b> 138:22 <b>impression</b> 29:19 31:15 47:4 48:24 49:25 <b>impressions</b> 20:25 28:7,10,10,11,11 28:19 29:4,21,22 31:15 47:12 51:7 51:8 65:11 <b>inaccurate</b> 166:24 <b>inbound</b> 69:18 73:2,24 74:4,7 127:22 128:1 <b>inception</b> 149:17 149:19 <b>include</b> 51:16 52:23 53:24 117:13,18 118:9 <b>included</b> 117:25 <b>includes</b> 56:19 123:16
---	--	--	---

## HIGHLY CONFIDENTIAL

[including - issuance]

Page 17

<b>including</b> 38:23 98:11 122:15,20 127:9 <b>inclusive</b> 59:16 64:14 <b>incoming</b> 130:2 <b>incomplete</b> 88:2 <b>index</b> 44:19 <b>indexed</b> 40:20 65:19,22 66:1 <b>indexes</b> 44:25 66:11 <b>indexing</b> 99:9 <b>indicated</b> 8:11,14 <b>indicates</b> 144:23 <b>indication</b> 92:21 <b>individual</b> 9:21 48:18 59:17 64:15 102:3 142:18 <b>individually</b> 141:18 142:19 <b>individuals</b> 32:11 48:21 49:10 86:13 <b>infer</b> 45:6 48:5 78:13 123:2 152:16 <b>inference</b> 118:17 <b>inferences</b> 78:4 93:10,17,23,25 96:25 118:21 119:5 124:15 125:4 132:15,19 <b>inferential</b> 121:16 121:16 <b>inferred</b> 6:3,6,7 46:8 56:19 57:6 59:16 62:13 64:14 65:2,4 78:16,19 88:9 89:1,13 91:25 92:13 119:22 122:23,24	123:5,11,14 124:24 125:1 150:12 <b>info</b> 98:9 109:8 <b>inform</b> 172:13,13 <b>information</b> 4:17 7:22 11:23 12:3 13:12,14 19:1 20:3 21:24,25 30:20,25 31:5 36:21,22 45:12 51:5,10,16 53:15 53:24 54:3 56:16 57:5 78:6 79:2 93:8,21 95:10 96:20 98:2 103:5 103:21 110:2 115:7,18,18 116:19,20 120:4 134:21 135:15,18 141:22 170:9,11 171:15 <b>informative</b> 171:1 <b>infrastructure</b> 33:11,22 60:18,20 61:17 77:12,15,16 84:21 87:13 91:18 91:19 127:24 152:22 <b>initial</b> 4:11 <b>inner</b> 171:23 <b>innovation</b> 103:18 <b>input</b> 122:24 123:1 <b>ins</b> 72:6 82:22 <b>inside</b> 10:22 44:22 45:1 54:12 61:7 81:8 147:7,8 <b>insight</b> 127:19 155:17	<b>instagram</b> 22:22 <b>installed</b> 11:10 12:10 13:13,15,16 106:25 109:7 <b>instance</b> 74:20 79:14 93:24 96:19 96:21 97:23 114:14 126:20 127:8 144:11 165:15 <b>instances</b> 23:15 79:18 96:11 97:4 97:10 129:9 165:20 <b>integer</b> 61:6 <b>integrate</b> 107:16 139:15 <b>integrated</b> 135:12 <b>integration</b> 10:6 20:1 106:5 125:10 125:13 134:17 139:24 140:2 <b>integrations</b> 9:15 141:5 <b>integrity</b> 37:8 158:7 <b>intentional</b> 51:13 <b>intentionally</b> 41:22 <b>interacted</b> 69:24 <b>interacting</b> 9:14 52:3 <b>interaction</b> 9:13 21:4 <b>interactions</b> 9:18 <b>interacts</b> 52:16 <b>interchangeably</b> 5:2 <b>interest</b> 8:14 101:19 166:14 172:11	<b>interested</b> 8:11 26:14 101:25 103:25 142:13,14 143:11 173:12 <b>interesting</b> 51:21 105:23 <b>interests</b> 139:19 <b>interface</b> 14:1,11 15:22 37:20,22 <b>interfaced</b> 76:20 <b>intermediate</b> 89:19,21 <b>internal</b> 36:5 56:8 57:9 58:9 90:1 125:6 157:25 159:6 171:16 <b>internally</b> 52:7 <b>internet</b> 67:23 <b>interplay</b> 158:25 <b>interrogatories</b> 24:10 119:25 <b>interrogatory</b> 139:11 141:17 <b>interrogatory's</b> 12:21 <b>intersected</b> 10:23 <b>intersection</b> 12:13 13:10 81:13 <b>intertwined</b> 71:14 <b>introduced</b> 149:9 <b>involve</b> 53:9 134:17 <b>involved</b> 77:24 138:19 162:11 166:6 <b>ios</b> 74:15 <b>ip</b> 23:7 26:1 39:17 39:19,20 42:2 51:6 152:7,8 <b>issuance</b> 170:20
---	---	--	--



## HIGHLY CONFIDENTIAL

[issue - learning]

Page 18

<b>issue</b> 5:14 62:15 70:13 146:15 171:6,20 <b>issues</b> 64:10 97:11 <b>issuing</b> 170:8 <b>item</b> 132:12 <b>iterated</b> 139:1	71:2 121:21 <b>keyed</b> 40:12 <b>keys</b> 66:25 67:2 <b>kick</b> 7:16 57:12 <b>kidding</b> 67:24 <b>killed</b> 33:6 <b>kind</b> 8:12,16 25:4 35:10 44:23,25 65:20 66:13 70:21 74:19 93:21 94:11 96:14 98:11,13 111:13 123:10 125:22 126:18 133:22 136:25 144:25 159:20 161:6 <b>kinds</b> 56:17 80:7 161:12 <b>knew</b> 32:9 81:21 <b>know</b> 6:19 16:14 16:15 17:4,6,14 21:18,21 22:1,5,11 23:2,5,12,14,15,23 24:17 26:6,12 28:8 30:13 31:24 32:1,22 33:8,9 34:3,6 35:9,22 36:8 39:4,24 40:9 42:9,10,16,20,22 43:1,13 44:9 47:5 47:13,25 48:5 49:10,19 50:1,4 51:7 54:20 55:19 56:10 58:13,14,16 60:23 65:8,10,11 66:15 69:9,13,22 70:3,6 71:5,18,21 71:22,25 72:4,5 73:3,15,16,23 74:14 75:10 76:8 76:18,22 78:15,21	79:12,17 82:1,22 83:7 85:21 88:12 89:3 91:1 92:1,16 94:12,14 97:1 98:8,19,21 100:9 100:11,15,16,17 102:11,15,23,24 106:6 107:22 110:10 115:15 117:18,25 118:9 118:11 119:12,13 119:20 120:5 121:21 125:1 126:9 130:11 132:18,20,21 134:23 135:10 136:7 138:15 140:17,19 141:1,3 141:7,10,16 144:17,18 145:5,7 147:1,1,2 149:3 153:9 154:5,8,9,11 154:16,19,24 155:4,10,20 156:8 156:15 157:4,5,20 158:14,15,17,22 159:17,18,25 160:23 161:16,25 162:4,23,24 163:15,17,23,25 164:2,4,15,24 166:16,23 167:14 168:3,6,8,9 169:21 170:1 171:8 <b>knowing</b> 46:13 76:6 156:20 <b>knowledge</b> 8:12 118:23 122:9 129:6 149:2 155:15	<b>knowledgeable</b> 8:15 49:9 86:15 104:8 159:20 <b>known</b> 137:19 <b>knows</b> 47:10 84:12 88:3 102:23 <b>kutscher</b> 3:16
<b>j</b>			<b>l</b>
<b>jams</b> 1:1,13 2:1,23 57:23 101:13 169:25 170:10 <b>jargon</b> 56:8 <b>javascript</b> 69:14 82:16,17 85:20 <b>je</b> 6:10 <b>job</b> 1:25 <b>jobs</b> 75:16 <b>joined</b> 84:4 <b>json</b> 147:7 148:5 <b>judge</b> 5:13 49:18 49:18,19 172:2,18 <b>jump</b> 146:7 <b>junk</b> 163:13			<b>labeled</b> 170:12 <b>laborious</b> 127:7 <b>lack</b> 13:7 <b>language</b> 22:8 40:20 41:2 <b>large</b> 32:6 63:6 122:10 147:5 156:1 <b>larger</b> 31:11 56:18 <b>largest</b> 147:11 148:11 <b>laser</b> 99:11 134:5 156:4 168:24 <b>latency</b> 134:1 146:15 <b>laufenberg</b> 3:5 <b>launch</b> 15:7 <b>laura</b> 3:17 <b>laurels</b> 121:3 <b>lawyer</b> 105:2 <b>lawyers</b> 6:20,23 6:24,25 160:15 172:14 <b>layer</b> 67:1 70:24 77:17 95:9,22 155:6 156:5 162:11 <b>layers</b> 116:23 <b>leaf</b> 168:7 <b>learn</b> 7:5,10 102:19 <b>learning</b> 102:1 161:17
<b>k</b>			
<b>kathryn</b> 2:25 67:25 <b>katy</b> 1:23 173:4,22 <b>keep</b> 18:2 20:22 26:15 32:8 35:13 51:15 83:12 92:10 93:23 125:9 <b>keeping</b> 101:20 <b>keller</b> 3:4 <b>kellerrohrback.c...</b> 3:8 <b>kept</b> 32:13 35:20 72:3 <b>key</b> 39:5,9,10 61:4 61:5 62:18 63:24 63:25 66:24 68:24			

## HIGHLY CONFIDENTIAL

[leave - machinations]

Page 19

<b>leave</b> 13:1 77:14 172:7 <b>left</b> 28:2,4 33:3 57:24 <b>legacy</b> 32:4,13,17 32:21 33:7 34:4 34:16 35:22 138:1 <b>legal</b> 50:5,7 <b>legs</b> 55:3 <b>length</b> 152:9,9 <b>lesley</b> 3:9 <b>lesser</b> 81:14 <b>letter</b> 27:9 43:10 169:25 <b>level</b> 15:15,16 18:16 42:19 46:23 48:22 78:18 85:2 86:24,24 97:3 110:24 115:16 121:15 123:9 160:25 <b>levels</b> 89:16 <b>leverage</b> 152:22 <b>leveraging</b> 87:25 123:10 <b>libraries</b> 18:15 <b>lifetime</b> 35:8 <b>light</b> 153:16 <b>likes</b> 51:8 <b>liking</b> 69:16 88:24 <b>limit</b> 15:23 158:12 <b>limited</b> 18:11 105:3 <b>limits</b> 128:11 <b>line</b> 102:6 103:14 125:12 148:17 <b>lines</b> 8:22 <b>linguistic</b> 90:13,17 90:19,23 <b>link</b> 21:14 68:3 105:15 109:20,24	110:3,5 116:12 <b>linked</b> 94:16 103:4 <b>list</b> 7:21 27:19 33:15,19 35:24,25 36:7 41:9,16 42:11 43:3,9,9,12 45:23 46:23 49:6 50:21 53:15 54:2 54:5 62:25 63:21 65:8 66:12,15 69:19 70:2 76:3,5 76:6 82:3 92:12 93:1 97:18,24 98:6 104:21 107:5 114:22 127:9 130:14 131:21 132:12 134:4 136:3,20,22 137:1 153:14,18 159:4 160:17 166:20 168:4,20 171:2 <b>listed</b> 57:24 165:7 165:20 <b>listening</b> 41:7 139:17 <b>litany</b> 52:16 65:6 <b>literally</b> 16:19 32:8 42:22 55:6 72:8 88:18 <b>litigation</b> 1:4 2:4 <b>little</b> 7:13 50:15 69:6,18 107:25 118:12 159:25 165:9 <b>live</b> 99:9 111:25 <b>lives</b> 34:1 38:2 <b>lldbpostmortem</b> 44:1 <b>llp</b> 3:4,9,15 <b>load</b> 28:8 31:10 122:11 134:7	166:3 <b>loaded</b> 69:15 82:17 85:10 <b>loading</b> 95:18 <b>local</b> 172:3 <b>location</b> 22:2 <b>lock</b> 121:21 <b>loeser</b> 3:5 103:10 <b>log</b> 14:2,3,11,23 15:2,4,8,14 16:6 17:17 18:4 21:23 58:4,5 59:4 62:3 79:11,16 80:12,16 80:18 107:10,19 108:2,17 112:13 134:20 <b>logdevice</b> 156:7 <b>logged</b> 29:17,20 47:14 52:20 79:12 106:25 107:4,7 <b>logging</b> 14:14 15:17 25:4 28:18 28:22 47:12,18,19 47:25 72:21 83:13 101:14,14 <b>logical</b> 88:5 <b>logs</b> 14:17 25:25 112:14 127:22 128:1 <b>long</b> 26:22 50:1 61:17 107:17 122:5 156:25 <b>longer</b> 105:20,22 <b>look</b> 4:18 7:13 11:11 24:21 26:22 40:15 41:1 42:5 42:21 43:15 45:12 52:6 63:12,22 65:17 69:17 84:1 84:7 86:22,22 91:12,12,13,15	93:14,14,15 95:4 97:7 105:20 127:4 127:11 128:6 130:12 133:13 145:4 153:14 <b>looked</b> 57:2 109:11 153:23 <b>looking</b> 27:9,19 28:1 37:14 46:23 49:13,14 60:23 82:3 85:11 98:13 127:21 128:1 131:19 165:25 166:2 <b>looks</b> 98:4 <b>loose</b> 85:18 <b>loosely</b> 70:14 140:19 <b>los</b> 4:1 <b>losing</b> 37:9 <b>lost</b> 67:20 166:11 <b>lot</b> 10:5 12:25 18:2 19:13 24:11,11 25:6,24 33:7 37:21 39:7 41:21 50:2,12 60:13,14 69:4 75:9 83:23 87:11 90:10 97:19 99:22 109:4 118:14 129:17 161:17 172:10 <b>lots</b> 6:11 33:8 34:20 152:10 <b>loud</b> 88:8 <b>love</b> 145:13 <b>lweaver</b> 3:12
			<b>m</b>
			<b>mac</b> 26:9 <b>machinations</b> 34:20

## HIGHLY CONFIDENTIAL

[machines - mean]

Page 20

<b>machines</b> 162:16 165:6 166:4,5	18:22 19:6 20:4 20:14 21:5,8,11,16	93:10,13 94:1,7,16 94:24 95:4,11,23	148:16,25 149:7 149:12,17,20,24
<b>mail</b> 39:25 68:6 76:14,18	22:7,12,21 23:4,19 24:4,6 25:8,11	96:6,16,23 97:5,15 98:4,14 99:1,18	150:4,9,16,19 151:1,16,22 152:4
<b>main</b> 82:17	27:8,15,21 28:3,16	100:4,8,14,20	152:8,15,25 153:6
<b>maintain</b> 41:21 87:23	28:24 30:1 31:13 31:21 32:2,24	101:9,16,19,23 102:6,16 103:14	153:12,20 154:2,7 154:10,13,18,22
<b>maintained</b> 41:23	33:5,13,18,24 34:8	104:12,15 105:1	154:25 155:9,11
<b>majority</b> 38:1 43:3 131:15	35:1,14,21 36:10 37:4,13,17 38:6,10	105:23 106:3,14 106:20 107:9,23	155:14,24 156:2,7 156:11,15,20
<b>making</b> 7:1 59:24 74:11 75:5 80:25 111:15 170:24	38:16 39:2,12,22 40:17 41:6,14 42:4,18,25 43:4,24	108:4,19 109:3,6 109:18,22 110:6 110:12,18 111:2	157:3,8,10,13,16 157:19,22 158:3,9 158:14,19,22
<b>manage</b> 36:16	44:5,12,14 45:4,17	111:16,21 112:19	159:3,23 160:5,11
<b>managed</b> 113:25	45:20 46:2,15	113:1,11,16,23	160:14,20 161:4,8
<b>manager</b> 104:18	47:3,9,15,21 48:2	114:2,9,21 115:5	161:13,19,23
<b>managing</b> 150:8	48:20 49:12 50:20	115:14,24 116:5	162:2,17,21
<b>manifold</b> 156:11 156:11,12 168:24	51:20 52:10 53:7 53:14,20 54:17,22	117:1,12 118:4,15 119:1,7 120:25	163:12,17,22 164:6,12,19,22
<b>manner</b> 40:3,5 58:23	54:25 55:5,25 57:11,18 58:1,7,11	121:12,23 122:3 122:12,17,23	165:1,4,8,22 166:7 166:13,18,25
<b>manually</b> 40:24	58:13,19 59:6	123:6,12,19 124:6	167:12,15,21,24
<b>mapped</b> 80:2	60:1,4 61:3,9,14	124:9 126:2,8,21	168:5,12,15 169:5
<b>mapping</b> 47:4 49:25	61:19 62:1,5 64:7 64:11 65:1,5,23	126:25 127:4,21 128:4,12,17,22	169:17,24 170:3,6 171:7,17,22 172:9 172:17
<b>maps</b> 48:24	66:2,7,14,23 67:3	129:1,4,10,16,25	
<b>market</b> 30:3 31:7 153:22 168:23	67:10,20 68:2,7,22 69:5,17 70:1,8,12	130:4,14,20 131:1 131:16,22 132:4,9	<b>master's</b> 8:8 159:21
<b>married</b> 78:14	70:15 71:3,9 72:4	132:13,18,24	<b>match</b> 32:23
<b>martie</b> 3:16	72:8,13,15,23 73:5	133:2,7,13,17	<b>matches</b> 66:9
<b>massive</b> 122:2	73:22 74:10,17	134:2,10,16	<b>matt</b> 3:10
<b>master</b> 1:13 2:23 4:4,6 7:20,24 8:2 8:4,6,10,13,17 9:10 10:3,16,18 11:14,21 12:5,11 12:15,20 13:20 14:8,24 15:10,13 15:20 16:2,7,10,12 16:19,22 17:10,13 17:21,25 18:9,13	75:1,13,20 76:7,22 77:1 78:3,10,25 79:9,21 80:19 81:2,9,15 82:5,14 82:25 83:5,11,18 84:1,5 85:1,14 86:16 87:6,8 88:1 88:16 89:11 90:1 90:3,9 91:10,23 92:2,5,18,24 93:7	135:22 136:2,9 137:2,7 138:4,8,20 138:25 139:9 140:8,11,21,24 141:2,7,13 142:9 142:13 143:14 144:4,12,21 145:15,24 146:4 146:17 147:10,16 147:22 148:7,13	<b>matter</b> 35:7 136:23 142:9 <b>mean</b> 7:8 14:5 19:13 21:23 22:3 24:10,11 25:21 30:17,23 32:25 33:6,8,24 34:9,15 37:20 39:22 42:8 42:19,20,21 43:11 44:6,10 45:16,25

## HIGHLY CONFIDENTIAL

[mean - narrow]

Page 21

46:20 50:10 51:14 51:14 52:4,11,12 52:12,13 64:16 65:5,6 68:22 70:19 71:9,17,18 73:22 78:10 80:1 83:6 84:2 85:16 86:21,21 88:5,5,14 88:16,18,25,25 90:1 92:5 97:15 97:17 114:13 119:8,18 121:13 128:20 129:13,16 129:23 130:9 144:21,25 145:16 146:17 152:13 157:22 159:15 166:22 167:15,21 167:23 171:14 <b>means</b> 11:1 31:15 31:16 52:7 73:4 139:8 <b>meant</b> 10:7 132:6 163:10 <b>measurement</b> 140:15,18 141:8 141:19 142:20 <b>mechanism</b> 60:6 62:6,16 <b>mechanisms</b> 149:11 <b>media</b> 34:12 <b>meet</b> 92:8 <b>meeting</b> 6:20 <b>melamed</b> 3:10 <b>mem</b> 95:7 <b>memcache</b> 93:2 95:8,22 99:5 157:4 163:8,11 168:25	<b>memory</b> 108:1 134:8 160:16 161:14 <b>mention</b> 137:5 <b>mentioned</b> 40:22 68:16 94:5 95:7 96:4 131:10 140:2 146:16 169:10 <b>mentioning</b> 63:11 124:4 <b>mentions</b> 40:25 <b>message</b> 91:8,12 91:13 93:14 110:16 121:6,6 131:16,20,24 140:4,6 <b>messages</b> 93:15 110:17 121:9 124:13 125:5 151:12 <b>messaging</b> 90:16 110:15 131:8,11 151:12 154:16,23 155:1 161:9,10 169:4 <b>messenger</b> 107:20 132:11 <b>meta</b> 151:14 <b>metadata</b> 111:20 152:10 <b>methods</b> 63:9 <b>mic</b> 101:15 <b>migrate</b> 33:12 35:15 <b>migrated</b> 35:15 38:20 <b>million</b> 29:10 <b>millions</b> 8:22 <b>mimics</b> 14:1 <b>mimmicks</b> 14:11	<b>mind</b> 5:8 6:2,14,16 7:16 14:7 24:14 59:13 65:4 73:21 79:20 87:20 88:2 <b>minimum</b> 37:1 84:13 100:9 124:23 <b>minute</b> 55:7,10,14 59:2 100:23 101:1 101:5 145:21,25 <b>minutes</b> 37:18 55:15,20,21 134:3 148:18 <b>missed</b> 24:12 80:14 <b>missing</b> 29:25 <b>misused</b> 8:10 <b>mm</b> 134:14 <b>mobile</b> 9:25 10:11 10:21,21 11:15,16 12:4,5,6,7 21:17 21:20,21 22:3,6,10 23:6,8,8,19,20,25 24:16 52:19 80:8 109:4,16 120:2 143:14,21,22,25 144:3,11,13,14,18 145:6,7,17 <b>model</b> 86:19 <b>modify</b> 40:15 86:4 127:13 <b>money</b> 50:2 164:10,17 <b>monitoring</b> 25:5 <b>months</b> 120:21 138:24,25 139:1 <b>morphed</b> 35:17 <b>motion</b> 171:19 172:2 <b>mouse</b> 21:8	<b>move</b> 9:1 32:11 57:21 101:20 171:5 <b>movie</b> 125:4 135:15 <b>movies</b> 135:1,2 139:18 <b>moving</b> 19:20 <b>multiple</b> 38:5,15 76:9,10 77:6 78:23 84:21 91:17 95:18 125:14 <b>multitude</b> 40:10 <b>mmm</b> 3:17 <b>music</b> 139:17 <b>mute</b> 8:18,18 <b>mysql</b> 38:6,9,17 38:22 43:5 66:23 67:5,6,15 93:3 94:2 96:11 158:21 162:2,5,5,6,7,10 162:12,13 163:9 165:4,6,7,9,12,12 165:20 168:25 <b>mystery</b> 118:15 <b>n</b> <b>name</b> 15:1 18:7 40:21,25 79:3 95:17 109:21 131:19 132:11 147:4 173:14 <b>named</b> 4:7,18,20 5:2,6 9:4 56:14,20 56:24 57:1,7 88:11,11 100:19 103:7 146:15 <b>names</b> 27:20 138:15 148:6 167:16 <b>narrow</b> 49:14 64:20 77:9
---	---	---	---

## HIGHLY CONFIDENTIAL

[narrowed - ones]

Page 22

<b>narrowed</b> 4:13 <b>narrower</b> 31:5 <b>naturally</b> 18:17 <b>nature</b> 171:11 <b>navigate</b> 82:2 <b>ne</b> 6:10 <b>nearby</b> 22:1 110:7 <b>nearly</b> 42:3 <b>necessarily</b> 52:8 77:18 128:15 <b>necessary</b> 7:2 35:9 41:23 45:23 <b>need</b> 19:12 24:7 27:6 58:4,4 62:3 71:6,13 78:8 82:1 84:16,22 85:5 86:21 107:3,4 120:18 143:15 146:25 147:2 151:17 155:7 159:8 160:3 169:1 <b>needed</b> 37:1 95:20 108:15 109:2 112:20 <b>needing</b> 106:25 137:20,21,21 <b>needs</b> 20:10 49:5 49:15 50:3 79:3 <b>neither</b> 154:12 156:10 157:12,15 157:18 162:1 163:16 173:10 <b>nesting</b> 13:7 <b>net</b> 84:23 <b>netflix</b> 74:12 79:6 79:14,25 82:6,9,14 84:7,15,16,23 86:8 118:6 125:3,16 134:25 135:2,8 139:21	<b>network</b> 21:22 128:7 144:24 <b>never</b> 63:23 76:12 <b>new</b> 32:14 33:23 40:24 45:18 86:4 93:21 131:18 151:10 <b>news</b> 117:8 131:12 <b>newspaper</b> 40:20 42:10 43:8,17 65:7,12 <b>newspapers</b> 40:19 <b>nine</b> 26:23 <b>nodding</b> 15:12 47:8 <b>nokia</b> 15:10 <b>non</b> 76:4 98:19 <b>nonmobile</b> 22:13 <b>nonpaid</b> 28:19 <b>nontechnical</b> 38:11 <b>nontrivial</b> 63:16 97:12 <b>nope</b> 136:1 157:14 164:21 <b>normally</b> 10:23 18:19 19:2 <b>note</b> 7:2 162:19 <b>noted</b> 162:21 <b>nother</b> 36:2 <b>notification</b> 74:15 168:20 <b>notifications</b> 74:16 161:3,4,7,9,10 <b>notified</b> 74:24 170:16 <b>noting</b> 34:2 142:2 <b>number</b> 18:19 27:20 29:21 33:1 37:7 40:1 41:3 42:15 61:3 86:2	120:21 148:1 166:5 <b>numbers</b> 23:6 61:11,18,20  <b>o</b>  <b>o0o</b> 1:2 2:2 3:22 4:3 172:22 173:3 <b>oakland</b> 3:11 <b>object</b> 6:24 7:3 52:11,16 60:5 64:23 93:22,24 <b>objective</b> 4:9 <b>objects</b> 61:16 94:21 148:6 <b>obligated</b> 109:1 <b>obtained</b> 5:23 6:5 46:6 53:16 62:11 108:20 119:21 124:19 <b>obvious</b> 30:7 <b>obviously</b> 23:5,22 48:7 68:20 104:9 <b>occurred</b> 29:19 <b>odd</b> 93:19 <b>offer</b> 139:13 <b>offerings</b> 23:15 <b>offhand</b> 34:7 42:17 71:21 78:22 <b>offline</b> 25:3 29:18 119:16 160:12,13 160:16 <b>oh</b> 9:5 20:14 24:6 32:24 33:18 37:17 40:16 43:5,5,24 57:18 59:3 67:23 68:14 101:2 108:6 133:7 134:2 163:7 163:7 <b>oil</b> 100:5 157:5 <b>okay</b> 4:6,25 5:15 5:20 6:13,16 7:7	7:20 8:3 9:5,10 10:16 11:21 12:23 13:20 14:12,15 15:20 16:10,21 17:13 19:17 20:4 22:7 23:4 25:16 25:18 26:16,19 31:21 32:2 34:8 37:17 38:10,12,16 39:2,12 40:1 42:4 42:25 44:14 47:3 54:17,24 55:22 58:7,9 61:21 62:1 68:2 69:5 75:20 82:25 85:15 90:7 90:9 94:7 98:14 99:1,6 104:23 107:9 108:4 110:6 112:19 128:5 131:22 132:13,24 133:2 134:16 138:8 148:13,16 148:25 149:12 154:22 157:3 158:19 161:13 162:17 163:19 164:22 165:1 166:7 169:15 171:21 <b>old</b> 26:9 63:18 107:14 <b>omnibus</b> 137:10 <b>omnistore</b> 157:6 <b>once</b> 27:1 29:12 45:24 69:6 116:14 169:21 <b>ones</b> 42:10 43:16 43:19 57:8,9 66:12,15 68:9 86:14 92:14,24 93:5 116:3,4
---	---	--	--



## HIGHLY CONFIDENTIAL

[ones - performance]

Page 23

119:9 127:9 136:4 156:21,21,22 159:10 164:4 <b>online</b> 67:12 163:14 <b>oozing</b> 99:6 <b>open</b> 15:2 86:3 115:17 <b>opens</b> 115:15 <b>operate</b> 16:16 66:13,15 107:15 <b>operates</b> 66:18,20 98:12 104:2 <b>operating</b> 26:4 <b>operational</b> 64:4 <b>operationally</b> 105:25 <b>operations</b> 36:15 <b>opportunity</b> 169:20 <b>opposed</b> 109:16 <b>order</b> 67:12 99:6 120:16 169:11 170:9,15,21 171:14 172:15 <b>orderdb</b> 163:24 <b>ordered</b> 92:9 <b>orders</b> 159:22 <b>ordinarily</b> 107:2 <b>ordinary</b> 72:19 79:19 167:19 <b>organization</b> 8:16 114:1 <b>organize</b> 55:13 <b>organized</b> 10:4,10 159:13 165:19 <b>original</b> 36:23 41:1 63:13 89:8 <b>originally</b> 63:23 <b>os</b> 52:20 152:8	<b>outcome</b> 173:12 <b>outline</b> 125:22 <b>outside</b> 22:19 60:14 79:18 91:15 106:2 136:8 161:8 <b>outward</b> 57:9 <b>overlap</b> 10:11 <b>overlayer</b> 98:7 <b>owned</b> 117:9 <b>owns</b> 11:2 98:24	142:12,14 144:4 144:14 146:18,19 159:1 163:9 168:8 172:2 <b>particular</b> 59:17 66:4 87:19 95:18 97:25 111:10,11 113:9 114:8 134:20 170:19 <b>particularity</b> 25:22 <b>particularly</b> 76:12 158:12 <b>parties</b> 5:23 6:5 11:23 13:2,5 46:6 50:8 62:11 79:5 79:19 124:8,20 139:24 170:24 <b>partition</b> 24:21 <b>partly</b> 131:10,11 160:13 <b>partner</b> 74:5,6,11 74:11 75:5 76:16 113:21 125:10,13 127:1,19 130:22 134:19,22 137:11 141:8 <b>partners</b> 9:15 10:7 75:14,15 125:15 128:7,19,23 136:5 136:8 139:13 140:15,18 141:19 142:20 <b>partnership's</b> 114:1 <b>parts</b> 46:14 47:17 47:18 <b>party</b> 9:23 10:20 13:25 14:9 17:17 18:4,16 19:7,25 76:1 81:17,20	106:5,7 108:25 112:23 119:21 123:17,20,23 125:6 170:14 173:11 <b>pass</b> 89:15 <b>passes</b> 116:14 <b>password</b> 15:2 18:8 19:4 <b>patent</b> 51:24 <b>patents</b> 51:22 78:3 78:11,14 91:1,5 121:4 156:3 <b>path</b> 20:5 <b>patient</b> 96:1 <b>patterns</b> 90:23 <b>pause</b> 59:5 101:18 170:5 <b>payer</b> 81:13,25 <b>payment</b> 43:7 <b>payments</b> 43:5,8 43:13,21 67:10,12 67:14,19 68:16,20 164:3,5,6,14,18 <b>pen</b> 39:23 <b>people</b> 13:13,14,18 29:4 32:14 34:1 37:8 45:7 47:24 56:9 58:15 60:10 61:10 78:12,23 88:19,22 95:25 104:10,21 135:12 138:19 142:11 160:2 166:1 168:7 171:13 <b>percent</b> 86:20,23 <b>perfect</b> 55:6,16 <b>perform</b> 42:1 <b>performance</b> 25:5 41:24 62:22 63:2
---	--	--	---

## HIGHLY CONFIDENTIAL

[period - present]

Page 24

<b>period</b> 85:2 125:18 149:2	126:17 151:13 156:18 164:13,16	<b>planet</b> 30:17	<b>position</b> 35:3 143:24
<b>periodic</b> 48:25	<b>phrase</b> 114:15	<b>platform</b> 5:19,22	<b>possible</b> 9:19
<b>periodically</b> 47:7	129:12	5:24 6:1,3,8,9,12	14:19,21 32:22
<b>permission</b> 11:11	<b>pick</b> 82:7	12:9 17:8 26:6	35:18 36:3 45:1
14:23 15:5 18:16	<b>pictures</b> 95:1	37:12 46:6,7,8	67:22 70:15 75:25
23:22 89:16	<b>piece</b> 51:24 60:5	48:12,12 50:13	77:13 91:7 97:3
107:18 109:19,23	77:15 84:20 87:12	51:13 62:10,12,14	98:15 111:1
<b>permissions</b> 12:10	111:14 145:9	69:8,11 72:16	127:24 145:6
12:19 106:17	<b>pieces</b> 50:3 77:12	74:3,20 75:5,21,24	147:10 148:2,3
108:11,17 112:13	77:16 91:17	79:5 119:20,22,23	170:21 171:5
137:21	152:22	123:13,22 124:1	<b>possibly</b> 27:3,16
<b>permitted</b> 15:1	<b>pii</b> 62:3	124:13,13,14,20	34:14 43:17 44:24
<b>person</b> 8:14 13:16	<b>pillar</b> 138:13,21	124:25 125:1	99:16 118:9
19:25 29:16 36:20	<b>pillars</b> 138:10	130:2 143:6	119:14
46:25 47:13 77:9	<b>ping</b> 58:8	<b>plausible</b> 136:21	<b>post</b> 53:13 54:12
117:8 143:11,12	<b>pixel</b> 56:23 70:9	<b>played</b> 9:19	54:16 61:1,1
<b>person's</b> 56:11	70:14,16,21,21	<b>please</b> 55:19	116:11 117:7
<b>personal</b> 93:5	75:21 77:2	<b>plug</b> 72:6,17,22	126:17 128:2,3
103:5	<b>pixels</b> 108:23	73:8 82:22	135:13 170:10
<b>personality</b> 78:5	<b>place</b> 9:13 65:13	<b>plus</b> 34:9,11,13,14	<b>posted</b> 51:17
90:24 93:17	120:19 121:1,3	34:17 70:16	52:24 53:25 56:18
<b>personalize</b> 57:10	173:7	<b>point</b> 17:4 20:11	81:25 117:9
<b>personalized</b>	<b>places</b> 85:8 125:15	24:22 33:10 41:16	131:13
56:13	135:16	44:25 55:6,16	<b>posting</b> 135:5
<b>personally</b> 66:19	<b>plaintiff</b> 4:18,21	72:3 76:13 90:11	<b>posts</b> 61:2 126:18
<b>perspective</b> 19:17	7:12 9:4 20:6	100:21 115:17	151:11
120:10	25:20 57:2	117:24 121:10	<b>potential</b> 4:14
<b>phone</b> 15:10 23:6	<b>plaintiff's</b> 4:7	122:4 135:8	<b>potentially</b> 4:12
23:17 24:1,1	68:23	141:17,24 142:1	4:14 28:22 40:14
39:25 61:3	<b>plaintiffs</b> 3:2 5:2,6	142:15 145:1	<b>power</b> 145:11
<b>photo</b> 53:12,12	7:3 9:7 16:23,24	<b>pointed</b> 44:22	<b>practical</b> 90:12
54:8,9,9,11,14,16	20:8 32:4 55:8,13	<b>points</b> 22:1 23:15	<b>pre</b> 109:7 113:23
63:15,21,22 81:24	55:25 56:15,24	<b>policies</b> 18:20 19:9	113:24 114:6
81:25 82:1,1,3	57:7 79:3 88:11	<b>policy</b> 19:14 35:7	137:2,9
94:10,14,15,21	100:25 101:24	37:10 136:23	<b>predict</b> 78:11
116:12,17,18,18	102:18 103:7	<b>poor</b> 61:17	<b>prefer</b> 145:12
151:14	146:7,10,15	<b>pope</b> 104:17	<b>presence</b> 164:23
<b>photos</b> 54:8 63:7	148:14 165:11	<b>populous</b> 137:13	<b>present</b> 14:1
63:18,20,24 82:3	169:13 170:18	<b>portal</b> 169:25	112:15 115:13
94:18,20,21	172:12,14	<b>poseidon</b> 100:11	
		157:10	

## HIGHLY CONFIDENTIAL

[presents - question]

Page 25

<p><b>presents</b> 14:10</p> <p><b>pretty</b> 38:3 46:20 50:23 61:10 98:5 104:7 119:6 162:10</p> <p><b>prevent</b> 15:22</p> <p><b>previous</b> 137:19</p> <p><b>previously</b> 165:18</p> <p><b>primarily</b> 65:22 80:23</p> <p><b>primary</b> 26:25 37:25 38:25 39:10 60:10 66:17 98:13 98:14 99:9 126:3</p> <p><b>prior</b> 97:20</p> <p><b>privacy</b> 1:3 2:3 62:3 89:15 103:19 104:4 126:17 160:1</p> <p><b>private</b> 80:2 106:8 113:2,6,12,15 114:3,12,17 115:7 117:14 118:5,6 122:7,15,20 124:4 124:22 128:6 129:11 139:12,23</p> <p><b>privilege</b> 6:20,21 7:4</p> <p><b>privileges</b> 81:14</p> <p><b>probabilistic</b> 85:12,15 86:19 151:4</p> <p><b>probabilistically</b> 92:17</p> <p><b>probability</b> 39:21</p> <p><b>probable</b> 39:19</p> <p><b>probably</b> 29:24 39:16 64:19 78:8 86:19 92:22 146:1 165:6 168:25 169:25</p>	<p><b>problem</b> 27:11,19 49:22 104:15 122:1</p> <p><b>problematic</b> 67:19</p> <p><b>procedures</b> 172:4</p> <p><b>proceedings</b> 1:12 55:24 59:5 101:8 101:18 146:3 170:5 172:21 173:6,11</p> <p><b>process</b> 33:15 40:2 40:4,10 83:24 103:12 111:25 113:20 121:14,20 136:24 137:4</p> <p><b>processed</b> 69:6 72:12</p> <p><b>processes</b> 37:11 152:11,14,17,20</p> <p><b>produce</b> 31:19 113:9</p> <p><b>produced</b> 73:17 102:21 103:11</p> <p><b>product</b> 25:2,10 25:11 27:6 35:11 35:16,18 47:24,25 59:23,24,25 60:17 71:7 74:20 96:21 117:6 121:9 131:6 168:8</p> <p><b>production</b> 25:7 33:3,16 110:25 133:15,18,19 158:1</p> <p><b>productive</b> 160:6</p> <p><b>products</b> 33:5 60:19 63:1</p> <p><b>profile</b> 1:4 2:4 45:15,24,24 51:17 51:19 52:1,5,7,24 52:25 53:2,4,22,25</p>	<p>54:6 56:9,10 57:9 64:16,18,19,22 90:15,19,20,23 91:14 95:19 124:14 152:17</p> <p><b>profiles</b> 51:22 93:16</p> <p><b>program</b> 85:20 115:8</p> <p><b>progress</b> 5:11 103:17</p> <p><b>project</b> 104:18 118:25 138:18</p> <p><b>properly</b> 38:14</p> <p><b>protective</b> 170:15 171:14</p> <p><b>protocol</b> 107:14 107:15</p> <p><b>provide</b> 4:8 27:1 56:12 103:10 106:17 118:20 122:7 143:16 169:24</p> <p><b>provided</b> 20:20 27:23 43:10 80:24 127:19 135:11 144:22</p> <p><b>provider</b> 108:21 145:17</p> <p><b>provides</b> 141:19 142:20</p> <p><b>providing</b> 141:25</p> <p><b>public</b> 7:22 107:2 126:16,17 127:13</p> <p><b>publicly</b> 106:9 113:7</p> <p><b>publish</b> 167:11</p> <p><b>published</b> 75:2 135:20</p> <p><b>publishing</b> 116:10</p>	<p><b>pull</b> 40:13 51:9 57:5 82:1</p> <p><b>pulled</b> 49:16 133:24</p> <p><b>pulling</b> 50:25 153:18</p> <p><b>pulls</b> 31:19,25</p> <p><b>purple</b> 88:17</p> <p><b>purposes</b> 4:20 70:11 162:14,15</p> <p><b>pursuant</b> 5:6 172:15</p> <p><b>pursuing</b> 139:18</p> <p><b>push</b> 74:14,15,20 74:21 161:9</p> <p><b>pushing</b> 126:24</p> <p><b>put</b> 9:2 28:9 30:8 53:5 68:7 145:25 159:8</p> <p><b>putting</b> 142:11</p> <p><b>pymk</b> 168:7,7</p>
			<p><b>q</b></p> <p><b>queried</b> 130:10</p> <p><b>queries</b> 41:4 44:21 89:14 132:14</p> <p><b>query</b> 36:18 81:20 81:22 88:21 95:19 96:24 97:3,6 102:5 103:13 126:20 129:20 130:6,11,12 131:14 146:13,19 146:23 147:2,21</p> <p><b>querying</b> 130:5,17</p> <p><b>question</b> 8:21,24 9:13 11:25 15:25 17:2,16 19:19 20:6 23:13 24:3,8 25:12 30:6,7 31:23 32:3 39:13 42:8,16 46:13,16</p>



## HIGHLY CONFIDENTIAL

[question - relatively]

Page 26

46:21 47:1,2 52:9 53:3,18 54:2 56:5 56:21,22 57:13,15 58:19,22 59:8,14 61:23 64:8,12 67:3 68:23 69:8 72:9 73:25 75:3 77:9 88:7,12,13 91:2,22 93:12 105:2 106:24 117:13,16 119:2 119:18 121:5 132:11 146:18,19 147:16,23 148:8 149:23 151:16 153:12 168:21 <b>questioning</b> 102:7 103:15 150:23 <b>questions</b> 6:17 7:18,21 9:7,11 16:24,25 19:23,24 20:9,11 28:4,6 50:6,7 55:9,13,18 55:21 56:1,6 58:25 100:23,24 100:25 101:11,24 103:16 104:10 106:15 140:14 142:10 143:25 145:10 146:7,10 148:14,19 159:17 159:21 160:3 169:13,18 <b>queues</b> 36:16 <b>quick</b> 58:6 <b>quickly</b> 160:7 <b>quite</b> 10:10 27:3 33:4 40:24 76:11 118:14 168:23 <b>quoi</b> 6:10	<b>r</b> <b>rained</b> 41:15 <b>raise</b> 155:16 171:8 171:9 <b>ram</b> 95:24,25 <b>randomized</b> 133:23 <b>randomly</b> 21:14 <b>rapidly</b> 153:5 <b>rare</b> 87:14 <b>rate</b> 78:13 128:11 158:11 <b>raw</b> 83:10,13 154:4 <b>read</b> 5:8,12 23:25 24:10,22,23 25:7,9 25:11 27:1,7 33:2 51:23 68:4 78:3 99:24 100:2 133:16 167:5 <b>readable</b> 118:1 <b>reading</b> 5:17 18:15 110:17 131:12,13 <b>reads</b> 116:4 <b>ready</b> 101:22 <b>really</b> 5:9 6:22 13:5 17:1 20:21 36:25 48:10 58:16 61:14 73:3 125:7 146:18 148:19 <b>realtime</b> 28:23 29:20 56:25 57:4 68:3 <b>reask</b> 57:13 64:11 147:22 <b>reason</b> 25:5 36:22 62:24 89:17 118:24 141:13 167:1	<b>reasonable</b> 4:18 5:5 9:3,4 118:12 <b>reasonably</b> 103:6 <b>reasons</b> 33:7 62:23 63:1 <b>recall</b> 54:7 72:11 108:1 109:10 110:8 121:10 128:1 131:18 135:9 140:1 156:18 157:1 158:24 163:8 <b>recalling</b> 74:14 <b>receive</b> 69:10,12 69:19 72:16 73:11 74:24 75:4 115:18 134:18 137:11 170:15 <b>received</b> 28:5 72:19 110:2,17 136:5 <b>reciprocity</b> 75:9 <b>reckon</b> 34:13 <b>recognize</b> 21:24 27:19 43:3,20 48:17,18 168:13 <b>recollection</b> 91:8 118:2 <b>recommend</b> 165:24 <b>reconnected</b> 65:16 <b>record</b> 4:5 55:11 55:14,23 58:12 59:7,10,13 91:7 101:6,7,17,21 113:2,4,10 146:2,5 146:9 158:9 162:20 169:2 170:4,7 171:18 172:19	<b>recorded</b> 119:15 <b>recruiting</b> 157:25 <b>red</b> 88:17 <b>refer</b> 7:25 56:9 59:15 60:2,7,9 64:13,17,18 71:24 72:2 <b>reference</b> 60:22 <b>referenced</b> 102:20 127:16 <b>references</b> 94:22 105:19 141:18 142:17 <b>referencing</b> 53:19 <b>referred</b> 24:25 <b>referring</b> 22:23 34:19 128:10 144:2 <b>refers</b> 139:23 <b>reflection</b> 135:23 <b>regarding</b> 5:24 46:7 62:11 119:21 124:20 <b>regards</b> 17:2 <b>relate</b> 20:21 150:13 <b>related</b> 59:24 88:23 103:13 109:12 110:15 140:9 173:11 <b>relates</b> 5:6 150:11 <b>relating</b> 79:2 119:13 <b>relational</b> 68:12 <b>relationship</b> 29:6 30:4 102:1 112:17 167:17 <b>relationships</b> 48:6 <b>relatively</b> 86:18 89:4
--	---	---	---

## HIGHLY CONFIDENTIAL

[relevant - right]

Page 27

<b>relevant</b> 36:18 88:10,15 89:2 112:12 114:15 125:17 126:19 149:2 <b>reliability</b> 163:6 <b>reliably</b> 97:13 <b>relief</b> 172:1 <b>reload</b> 28:8 <b>relying</b> 108:20 <b>remain</b> 16:6 92:19 <b>remainder</b> 159:4 <b>remaining</b> 27:12 77:4,4 <b>remember</b> 6:19 34:9 46:3 82:18 83:7 107:21 108:6 131:19 156:13,17 163:4 <b>remittances</b> 164:17 <b>remotely</b> 75:18 <b>remove</b> 42:6 91:14 160:16 <b>removed</b> 17:7 137:14 159:5 168:19 <b>renamed</b> 162:25 <b>render</b> 95:21 108:14 110:4 <b>rendered</b> 53:4 <b>rendering</b> 113:6 <b>repeat</b> 64:10 93:9 93:11 109:21 149:22 <b>repeatedly</b> 31:22 95:10 <b>replacement</b> 58:14 102:1,7,10,19,25 104:24 112:24	<b>replication</b> 38:13 <b>report</b> 71:25 <b>reported</b> 1:22 <b>reporter</b> 68:1,5 80:13 101:22 173:5 <b>reporter's</b> 1:12 <b>repositories</b> 4:10 4:11 <b>repository</b> 26:25 27:5 127:16 167:4 167:6 <b>represent</b> 69:2 <b>represented</b> 141:4 <b>request</b> 42:1 80:24 81:1 83:10,14 95:9 102:3 103:20 109:19,23 110:5 111:15 128:2 170:10,11 172:11 172:12 <b>requested</b> 114:18 170:9 <b>requesting</b> 170:17 <b>requests</b> 22:18 79:16 80:5,10 109:15 110:2,3 128:3,11,14,15 130:2 135:13 158:10,12 169:9 <b>require</b> 7:13 18:19 40:23 86:7 89:16 <b>required</b> 87:1 106:6 <b>requirements</b> 103:20 <b>requires</b> 49:24 <b>research</b> 44:6,7 133:20 166:14,15 166:19 167:4,5,6 167:10,20	<b>reset</b> 101:15 <b>resolve</b> 171:6,8,10 <b>resource</b> 84:8 <b>resources</b> 104:16 104:19 128:16 <b>respective</b> 172:4 <b>response</b> 12:21 114:25 <b>responsible</b> 47:11 49:9 85:6 86:9,14 87:24 105:6 <b>rest</b> 121:3 136:14 137:19 138:1 <b>restricted</b> 17:7 109:14 <b>rests</b> 96:17 <b>result</b> 8:25 58:1 79:15 80:11,15,18 84:23 111:13 112:7 136:25 <b>results</b> 72:20 <b>resume</b> 55:15,21 101:6,9 <b>retain</b> 35:8 105:10 <b>retention</b> 48:1 <b>rethink</b> 20:6 <b>retrieve</b> 11:3 70:20 90:15,18 <b>retrieved</b> 130:23 <b>return</b> 12:9,12 69:6 82:5 122:21 123:4 <b>returned</b> 12:17 114:18 125:3 <b>returns</b> 66:8 <b>reverse</b> 54:16 <b>review</b> 18:19 111:25 113:20,21 120:3 121:4 136:24 137:4	<b>reviewed</b> 106:7 <b>rfalconer</b> 3:20 <b>rhythm</b> 30:24 31:3 <b>richer</b> 49:4 <b>rid</b> 102:14 103:1,2 104:23,23 105:13 105:15,21,22,25 106:2 168:25 <b>rids</b> 105:7 <b>right</b> 4:4,4 6:22 9:15 13:1,20 15:18 16:8,9 17:3 17:3,5 18:13 19:6 19:18,25 20:17,17 21:1,2 22:2,3,12 22:24 23:9,10,20 24:14 25:12,20,23 26:1,8 28:10,11,25 30:2,21,23 31:14 33:6,7 34:10,16,25 35:21 37:4,20 38:7 39:24 41:10 42:7 44:15 45:8 45:11,18 46:20 47:4,5,7 48:15,23 48:24 49:20 51:5 51:14 52:2,18 55:16 58:6,17 62:8,10,20 64:16 65:11 66:8 69:5 70:8 71:12 73:21 73:24 74:3 75:2 79:22 85:14 88:7 90:7,15,19,25 91:11 96:18 97:18 97:22 99:20 100:25 101:9,20 101:23 103:20 104:10 106:3,12 111:21 112:3,3 115:19,20 118:10
--	--	---	---

## HIGHLY CONFIDENTIAL

[right - server]

Page 28

119:18 121:23 122:25 124:11,12 124:21 125:10 126:10,11 130:4 130:12 132:9 134:20 137:7 140:12,13,25 143:10,12 145:15 146:4 147:12 148:2,8 152:4 153:19 154:3,9 157:3 159:10 160:2,14,19 161:23 163:12,14 164:19 165:13,14 168:22 169:7,15 169:19 170:3,6,8 <b>ring</b> 3:18 167:14 167:23 171:7,21 172:6 <b>rmr</b> 1:24 2:25 173:22 <b>rohrback</b> 3:4 <b>room</b> 146:1 <b>rose</b> 3:18 <b>ross</b> 3:19 <b>roughly</b> 149:6 <b>rpr</b> 1:24 2:25 173:22 <b>rsc</b> 157:13 <b>rudimentary</b> 41:12 <b>rule</b> 87:5 99:3,16 172:5 <b>rules</b> 172:3 <b>ruling</b> 171:20 <b>run</b> 34:17 61:10 76:21 160:7 162:13 164:14 <b>running</b> 10:20 61:18 77:11,14	154:1 164:10 <b>runs</b> 164:18 <b>russ</b> 3:16 73:16 117:17 121:17 138:7 <b>russian</b> 13:7 <b>s</b> <b>s</b> 133:4 <b>s3</b> 157:2 <b>sais</b> 6:10 <b>sample</b> 86:18 97:16,16 <b>sandeep</b> 58:8 59:3 <b>sat</b> 90:10 <b>satisfaction</b> 95:13 <b>satisfy</b> 102:2 <b>save</b> 73:23 <b>saw</b> 30:5 73:24 109:3 120:12 <b>saying</b> 6:24 19:9 43:9 65:14 68:23 75:15,15 86:6 142:4 <b>says</b> 48:11 62:9 83:20 125:1 141:16 <b>scale</b> 31:12 67:9 94:6 166:2 <b>scenario</b> 12:4,7 17:17 18:5 105:9 <b>scenarios</b> 35:10 <b>sces</b> 142:23 <b>scheduled</b> 129:1 <b>scheduling</b> 170:11 <b>schema</b> 52:13 147:19 <b>schemas</b> 147:6,12 148:12 <b>schmidt</b> 1:23 2:25 173:4,22	<b>scope</b> 56:18 138:18 159:1 <b>scoping</b> 120:23 <b>screen</b> 27:10 43:16 <b>script</b> 97:6,7,9 102:5 <b>scroll</b> 26:23 <b>sdk</b> 69:14,18,23,24 70:14 74:9 82:16 82:17 83:7,11 84:11 85:18,19,20 86:2,4 <b>se</b> 36:9 <b>seal</b> 171:14 <b>sealed</b> 171:13 <b>seamless</b> 139:14 <b>search</b> 33:25 34:2 44:16,20,25 62:7 126:7,11,14,19 155:19 <b>searches</b> 44:23 <b>searching</b> 27:22 44:17 <b>seat</b> 151:11 <b>seattle</b> 3:7 <b>second</b> 27:10 46:16 62:11 125:7 146:18 170:4 <b>section</b> 16:24 57:12 58:21,24 111:9 <b>secure</b> 80:22 <b>secured</b> 9:23 <b>security</b> 26:7 <b>seduce</b> 86:17 <b>see</b> 11:5 13:11,12 13:12 18:24 27:23 29:8 36:17 48:2 50:16 57:24 58:5 58:10 64:3 84:7 101:15 110:1	126:18 127:11,12 127:17 129:17 145:5 148:17 168:19 <b>seeing</b> 62:24 72:12 72:22 <b>seeking</b> 56:12 139:14 <b>seen</b> 105:7 157:24 167:18 <b>sees</b> 36:14 <b>segments</b> 57:6 <b>select</b> 42:6,6 139:13 <b>semantics</b> 137:20 137:22 <b>send</b> 68:3 82:24 140:4,6 164:17 <b>sending</b> 71:23 110:16 164:10 <b>sense</b> 12:21 13:21 34:24 37:5 40:1 45:8 54:18 133:8 148:6 167:16 <b>sent</b> 22:18 72:19 <b>sentence</b> 129:14 <b>separate</b> 91:9 131:9 144:10 155:5 <b>separately</b> 121:17 <b>series</b> 4:8 <b>serve</b> 33:16 41:4 44:23 63:5 133:23 159:7 166:1 <b>served</b> 82:17 142:10 <b>server</b> 9:22,22 10:12,12 11:9,9,18 70:24 74:24 80:10 80:10 81:6 157:19
--	---	--	---

## HIGHLY CONFIDENTIAL

[servers - speaking]

Page 29

<b>servers</b> 9:24,24 10:1 11:20 30:17 67:7 71:23 83:10 107:16,19 108:2,3 109:16 129:5 145:17 <b>serves</b> 99:12 <b>service</b> 155:20 156:13 158:6 168:10 <b>services</b> 52:16 79:11 144:22 164:13,23 <b>serving</b> 76:8 <b>session</b> 57:24 80:6 <b>set</b> 6:14,16 17:6 25:1 28:4,5,18 36:4 47:24 59:15 64:13 67:1 76:17 77:18 85:17 97:21 97:22 104:1 107:15 131:5 135:1,2 136:17 137:19 139:23 140:24 143:15 145:10 147:5 153:7 158:21 165:6 169:9 173:7 <b>sets</b> 28:23 31:5 76:17 111:18 <b>settings</b> 22:16 <b>sever</b> 105:18 <b>sftp</b> 129:5 <b>shard</b> 97:14 <b>sharded</b> 158:21,21 <b>share</b> 27:10 43:11 43:16 60:20 61:15 71:18 72:17 73:8 73:10 74:7 75:21 82:23 123:17	<b>shared</b> 123:23,25 126:16 <b>shares</b> 135:2 <b>sharing</b> 139:18 140:3 <b>shed</b> 153:16 <b>shiny</b> 33:23 <b>short</b> 99:6 <b>shorthand</b> 173:4,8 <b>show</b> 26:10 27:10 29:9,12,22 109:2,2 126:19 <b>shown</b> 29:15,16 <b>shut</b> 33:16 35:11 <b>sic</b> 29:2 <b>side</b> 11:18 22:4,6 22:10,13 23:8,25 69:23 75:19 77:7 101:10 108:20 112:15,16 <b>sides</b> 6:20 <b>sift</b> 120:22 122:2 166:6 <b>sign</b> 170:15 171:13 <b>signal</b> 21:22 120:1 <b>signature</b> 173:21 <b>significant</b> 40:14 42:15 47:1 <b>significantly</b> 86:18 <b>similar</b> 28:14 37:8 37:11 58:25 74:19 89:4 109:12 122:1 140:22 163:5 <b>simple</b> 20:14 46:4 46:20,21 60:16 118:25 <b>simpler</b> 77:22 <b>simply</b> 100:1 <b>single</b> 40:10 45:24 46:25 48:11,18 77:9 97:13 137:9	<b>sit</b> 51:6,8 <b>site</b> 19:3,4 21:4 135:11,11,12 167:10 <b>sits</b> 46:9 131:25 135:3 <b>sitting</b> 45:4 48:9 51:2 94:18 <b>situation</b> 79:1 <b>six</b> 138:24,25 <b>size</b> 120:21 148:11 <b>sizes</b> 31:6 <b>skills</b> 8:18 <b>skip</b> 97:16 <b>skipped</b> 68:2 <b>slash</b> 10:12 <b>slightly</b> 103:10 113:8 136:15 156:24 <b>slow</b> 134:6 <b>slower</b> 134:9 156:6 <b>small</b> 31:10 33:1 136:17 <b>smaller</b> 160:18 <b>social</b> 11:22 12:23 12:25 13:1,5 25:1 25:15 34:11 52:13 72:6 82:22 121:7 121:8 122:8,15,22 144:15 149:14,20 149:25 150:3,11 150:14 168:17 <b>socialized</b> 92:2 <b>software</b> 40:24 95:24 162:14 <b>solanki</b> 58:8 59:3 <b>solution</b> 120:13 <b>somebody</b> 44:25 56:18 69:15 74:12 116:9 118:6	<b>someone's</b> 78:13 <b>somewhat</b> 77:22 83:21 <b>soon</b> 24:24 <b>sorry</b> 9:6 17:8,19 17:21,22 19:19 40:8,16 41:6,7 43:24 50:9 55:1 55:11,12,17 64:9 70:18 80:13 81:8 83:12 88:8 89:23 91:10,14 130:21 132:5 133:9 138:11 <b>sort</b> 5:1 9:18,18 13:6 15:21 19:10 36:5 37:6 58:21 58:24 59:14 72:21 94:24 96:17 106:4 145:2 169:11 <b>sorts</b> 78:11 <b>sound</b> 92:9 99:19 119:10 124:22 <b>sounds</b> 71:15 75:22 94:1 96:25 99:19 113:18 115:25 120:25 121:1 128:5 139:22 143:15,23 145:16 <b>source</b> 8:22 69:15 72:21 167:18 <b>sources</b> 149:13,25 150:5,20 159:18 159:24 <b>space</b> 61:16 <b>speak</b> 28:12 73:19 94:25 163:1 166:21 <b>speaking</b> 54:13
---	---	--	--

## HIGHLY CONFIDENTIAL

[special - square]

Page 30

<b>special</b> 1:13 2:23 4:4,6 7:20,24 8:2 8:4,6,8,10,13,17 9:10 10:3,16,18 11:14,21 12:5,11 12:15,20 13:20 14:8,24 15:10,13 15:20 16:2,6,7,10 16:12,19,22 17:10 17:13,21,25 18:9 18:13,22 19:6 20:4,14 21:5,8,11 21:16 22:7,12,21 23:4,19 24:4,6 25:8,11 27:8,15,21 28:3,16,24 30:1 31:13,21 32:2,24 33:5,13,18,24 34:8 35:1,14,21 36:10 37:4,13,17 38:6,10 38:16 39:2,12,22 40:17 41:6,14 42:4,18,25 43:4,24 44:5,12,14 45:4,17 45:20 46:2,15 47:3,9,15,21 48:2 48:20 49:12 50:20 51:20 52:10 53:7 53:14,20 54:17,22 54:25 55:5,25 57:11,18 58:1,7,11 58:13,19 59:6 60:1,4 61:3,9,14 61:19 62:1,5 64:7 64:11 65:1,5,23 66:2,7,14,23 67:3 67:10,20 68:2,7,22 69:5,17 70:1,8,12 70:15 71:3,9 72:4 72:8,13,15,23 73:5 73:22 74:10,17	75:1,13,20 76:7,22 77:1 78:3,10,25 79:9,21 80:19 81:2,9,15 82:5,14 82:25 83:5,11,18 84:1,5 85:1,14 86:16 87:6,8 88:1 88:16 89:11 90:1 90:3,9 91:10,23 92:2,5,18,24 93:7 93:10,13 94:1,7,16 94:24 95:4,11,23 96:6,16,23 97:5,15 98:4,14 99:1,18 100:4,8,14,20 101:9,16,19,23 102:6,16 103:14 104:7,12,15 105:1 105:23 106:3,14 106:20 107:9,23 108:4,19 109:3,6 109:18,22 110:6 110:12,18 111:2 111:16,21 112:19 113:1,11,16,23 114:2,9,21 115:5 115:14,24 116:5 117:1,12 118:4,15 119:1,7 120:25 121:12,23 122:3 122:12,17,23 123:6,12,19 124:6 124:9 126:2,8,21 126:25 127:4,21 128:4,12,17,22 129:1,4,10,16,25 130:4,14,20 131:1 131:16,22 132:4,9 132:13,18,24 133:2,7,13,17 134:2,10,16	135:22 136:2,9 137:2,7 138:4,8,20 138:25 139:9 140:1,8,11,21,24 141:2,7,13 142:9 142:13 143:14 144:4,12,21 145:15,24 146:4 146:17 147:10,16 147:22 148:7,13 148:16,25 149:7 149:12,17,20,24 150:4,9,16,19 151:1,16,22 152:4 152:8,15,25 153:6 153:12,20 154:2,7 154:10,13,18,22 154:25 155:9,11 155:14,24 156:2,7 156:11,15,20 157:3,8,10,13,16 157:19,22 158:3,9 158:14,19,22 159:3,21,23 160:5 160:11,14,20 161:4,8,13,19,23 162:2,17,21 163:12,17,22 164:6,12,19,22 165:1,4,8,22 166:7 166:13,18,25 167:12,15,21,24 168:5,12,15 169:5 169:17,24 170:3,6 171:7,17,22 172:9 172:17 <b>specialize</b> 34:1 <b>specialized</b> 63:7 64:3 139:3,5,10 142:24	<b>specific</b> 4:23 13:9 14:22 20:8 23:14 27:3 28:15,21 31:10 32:17 36:16 36:17,18 44:20 47:25 56:6,22 59:24 63:4 64:5 69:20 71:10 75:15 77:23 80:9 82:6 84:19 85:6 87:25 91:21 96:21 98:24 106:15 111:19,23 112:10 116:15,24 125:20 126:1 129:14 131:5,6 140:7 141:14 146:23 150:9 151:8,25 152:21 152:21 153:3 161:12 165:21 168:17 169:9 <b>specifically</b> 8:9 9:14 23:24 103:24 141:3 143:20 168:6 <b>specification</b> 63:13 <b>specificity</b> 20:12 <b>specified</b> 10:12 <b>specify</b> 108:14 <b>speed</b> 23:7 <b>spelled</b> 108:24 <b>spend</b> 55:20 <b>spent</b> 22:10 <b>spot</b> 8:24 <b>spotify</b> 139:21 <b>springer</b> 3:6 102:22 <b>sql</b> 115:15 <b>square</b> 98:22
---	---	--	---



## HIGHLY CONFIDENTIAL

[stability - substantially]

Page 31

<b>stability</b> 25:5	80:4,17,21 82:20	<b>stolen</b> 122:6	112:18 116:16
<b>stack</b> 84:20 115:22	83:1,3 86:1 87:7	<b>stopped</b> 15:21	118:18,19 120:9
<b>stand</b> 100:9 168:7	87:18 92:21 93:1	<b>storage</b> 25:4 28:15	123:14 134:23
170:20	93:9,11,19 94:4,9	37:25 38:2,4,25	136:4 141:20
<b>standard</b> 75:2	94:19 95:8,15	63:6,9 67:12	142:21 144:8,19
109:8 130:11	96:5,8,19 97:2,24	77:17 91:9 94:11	156:16
<b>stands</b> 100:17	98:8 100:7,13,16	100:14 116:18,22	<b>stores</b> 45:18 88:18
<b>start</b> 6:17 53:2,23	102:13,14 104:25	116:24 128:8	96:20 128:23
60:15,16 65:13	105:3 106:1,11,12	151:10 154:4	131:25 153:24
77:20,22 79:24	106:19,22 107:11	155:5,7 156:3,6,13	<b>stories</b> 131:13
121:1 127:8	108:8 109:10,21	156:24 157:13	<b>storing</b> 30:24
153:20	109:25 110:8,14	158:16 160:11	31:10 84:9 99:5
<b>started</b> 62:25	110:21 111:8	163:14 164:23	133:11 156:1
137:5 149:5	112:2,3,11 113:3,4	165:18	165:11 166:3
<b>starting</b> 89:6,8	113:13,19,24	<b>store</b> 25:1 26:13	<b>story</b> 117:8 135:19
142:7	114:5,14,23	29:13,14 30:17	<b>strange</b> 83:22
<b>state</b> 16:1 103:6	122:12,13,19	31:7 38:20 39:3	<b>strategic</b> 125:15
141:24 173:1	123:3,8,15 124:3,7	39:13,17 40:11,12	126:2 128:6,18
<b>stated</b> 119:24	125:19 126:5,13	41:24 45:12,23	<b>strategy</b> 20:7
142:21 168:17	126:23 127:2,6,23	54:8 65:16 67:8	<b>stream</b> 115:17
<b>statement</b> 131:3	128:10,13 129:8	78:24 90:23 91:3	116:10
<b>statements</b> 167:3	129:12 131:3	91:17,19 93:7,18	<b>street</b> 3:11
<b>statistically</b> 86:17	133:6 135:9 136:1	93:20 94:9 96:7	<b>stretch</b> 55:3
<b>statistics</b> 138:15	136:13,17 137:4	106:2 120:11	<b>strikes</b> 167:6
<b>statute</b> 103:18	137:12,13,24	121:9 124:15	<b>stringent</b> 103:20
<b>stays</b> 122:6	140:22 141:3,10	128:2 132:15	<b>structure</b> 114:24
<b>step</b> 84:22,25 85:5	144:2,6 145:14,23	134:5 135:15,18	<b>struggle</b> 120:16
<b>steps</b> 113:22	148:23 149:3,8,15	135:19 144:10	<b>struggling</b> 71:12
<b>steven</b> 15:24 16:2	149:19,22 150:1,6	151:22 155:25	<b>stuff</b> 30:17
16:4,9 17:10,11,24	150:13,18,22	157:2 158:25	<b>sub</b> 12:13
23:12,12,14,23	153:10,11 154:9	159:8 169:4	<b>subject</b> 142:9
31:24 32:1,19	154:12,21 155:12	<b>stored</b> 22:15 24:15	<b>submission</b> 43:11
35:5,16 42:25	155:22 156:9,10	24:18 25:20 29:8	<b>submit</b> 24:4,8
43:2,16,18 44:4	157:8,9,12,15,18	30:4,14,16,20,21	<b>submitted</b> 117:23
46:19 48:15,16	158:18 160:10,25	38:18 40:3,5	<b>subscribed</b> 75:2
49:7 57:21 58:8	161:6,11,21 162:1	44:18 47:7,9,10	173:13
64:7,9,18 65:3,14	162:24 163:5,16	49:1 54:10,11,12	<b>subsequently</b> 4:13
65:24 66:3,10,17	163:21 164:4,20	66:6 69:4 71:1,1	<b>subset</b> 10:25 20:3
70:2,6,10 72:5,10	164:21 165:2,3,17	72:14 78:20,22	25:15
72:14,18 74:19	166:12,17 168:6	82:10 83:6,16,20	<b>substantially</b>
75:12 76:25 79:10	168:13	85:24 98:23	160:18 171:3

## HIGHLY CONFIDENTIAL

[succinct - tao]

Page 32

<b>succinct</b> 59:10 <b>sucked</b> 33:6 <b>sufficient</b> 39:18 62:22 <b>suggest</b> 104:6 168:9 <b>suite</b> 3:7,11,19 <b>supervision</b> 173:9 <b>supplemental</b> 43:10 <b>support</b> 36:11 67:18 68:12,18,19 94:6 171:18 <b>suppose</b> 40:19 <b>supposed</b> 84:12 145:25 <b>sure</b> 16:5 17:12 28:9 29:11,13,14 29:20 31:18 34:25 35:2 52:9 57:16 61:10 68:5 72:2 76:15 80:13 89:24 106:1 108:1 109:14 123:3,11 124:16 127:6 130:8 132:2,7,7 135:9 139:7 143:8 150:22 158:4 161:6,11 168:23 172:13 <b>surfaces</b> 60:19 <b>surprised</b> 77:8 <b>surveyed</b> 104:20 <b>synapse</b> 157:16 <b>synonymously</b> 96:10 <b>synthetic</b> 81:23 <b>system</b> 11:24,25 12:1 16:16 26:4 28:20 29:20 30:23 33:2,3 34:11,16	38:1 40:13,15 41:25 42:1,21,22 44:19,20,21 48:4 48:11 49:3,11 52:3,15,19 54:7 63:12,15,17,20,22 63:25 64:25 66:7 66:10,17,18 77:21 77:21 78:1,2 85:21 87:22 89:4 96:13,22,24 99:9 99:11,17,20 100:2 100:19 105:24 121:11 123:10 131:4 132:14,20 133:25 134:4,9 155:6,7 160:13,13 162:25 <b>systems</b> 4:14 5:5 7:6,11,14 13:24 24:20 25:7,13,17 25:19 27:3,12,16 28:14,14,15,21 29:25 30:2,6,12,16 30:21 31:1,4,11 32:4,7,8,10,12,13 32:18,22 33:11,19 34:4 35:15,23 37:6,8,11 40:4,4 40:11 41:3 42:5,5 42:11,19 43:12,14 44:18 45:2,10,12 45:22 46:11 48:3 48:9 49:6,8 50:22 50:25 51:1 52:17 60:22 61:7,22,25 62:7,16,21,22,24 63:3,4 64:3 65:23 66:11 68:25 69:4 70:17,21 71:1,2,13 73:14 77:3,4,5,25	78:2,24 81:18 85:24 87:8,11,15 87:19,24 91:24 92:9,14 98:5,9,17 98:20 99:3,14,16 99:24 100:1,10 105:25 106:1 110:19,20 119:13 120:8,10,22 121:2 122:10 124:16,17 125:7 131:20 132:19 134:6,12 134:12,12 136:4 141:15 143:4,19 150:25 151:7,8 152:23 153:5,15 153:16 155:16 159:5,6,6,7,7,8,9 160:7,16,17 161:17 162:5,6 165:11,18,23,25 166:2 168:18,22 170:18,19 171:2 171:16,24,25  <b>t</b>  <b>table</b> 99:13 104:22 130:11 147:8 <b>tableau</b> 30:16,22 157:19 <b>tables</b> 29:17 89:19 89:21 130:10,18 146:24,24,25 147:1,2,3,6,11,13 147:24,25 148:1 148:11 <b>taco</b> 162:25 163:4 163:8 <b>tacos</b> 167:21,22 <b>tacosrs</b> 162:22 <b>tagged</b> 54:9	<b>tags</b> 53:11 <b>tailer</b> 168:7 <b>take</b> 6:9 36:12 37:11 40:14 42:11 42:14 47:1 50:1 50:14 54:25 55:7 62:3 72:25 73:1 73:11 77:2,6 78:22 85:13 100:21,23 101:5 120:24 125:11 145:11,20,20,24 146:6 165:1 <b>taken</b> 36:20,23 55:24 101:8 146:3 173:6,7 <b>takes</b> 66:8,10 72:24 97:12 <b>talk</b> 12:25 13:23 22:13 37:18 59:20 79:21 82:6 84:8 84:15 104:16 106:4 119:16 141:8 143:12 168:1 <b>talked</b> 6:23 24:16 53:20 80:7 113:20 160:20 <b>talking</b> 12:3 14:4 43:14 46:2 86:21 114:3 118:4,5 <b>talks</b> 133:14 150:11 <b>tally</b> 158:5 169:6 <b>tangible</b> 9:2 42:20 50:3 <b>tao</b> 37:16,17,18,20 37:24,25 38:2,3,3 38:16,20 39:1 44:22,22 45:19,21 54:5 64:23 66:18
--	---	---	---

## HIGHLY CONFIDENTIAL

[tao - three]

Page 33

67:1,7,15,17,19 68:9,11,12,19,21 93:2,7,19,21 94:20 94:22,24 95:5,19 98:3,4,9,12 99:21 106:2 110:22 111:4 112:5,18 116:20 121:1,1 124:17 130:25 131:10,14,15 134:7 135:16,18 151:5 155:2,3 158:14,25 162:7 162:10 163:5,8 169:6 <b>taos</b> 98:6 <b>target</b> 76:5,14 <b>targeted</b> 57:10 <b>targeting</b> 76:2,16 76:19 78:9 123:10 123:16 132:22,25 133:1,8 <b>task</b> 28:15 36:16 98:23 115:21 <b>team</b> 47:1,12 49:8 78:8,9,23 84:18 85:17 91:17,21 96:20 98:24 104:4 104:7 105:5 120:21,22 132:23 132:25 133:8 138:7,9,10,21,23 141:8,12 <b>teams</b> 36:16 49:11 77:7 84:19,21 85:6 86:8,11,13 87:20,23 158:7 159:7 <b>tech</b> 67:17 <b>technical</b> 4:17 8:20 32:7 38:11	50:6 64:9 97:3 104:18 106:15 110:24 <b>technically</b> 14:20 110:25 <b>technology</b> 171:24 <b>telco</b> 109:4 <b>tell</b> 48:4 49:24 51:20 92:14 98:16 120:3 144:22 145:17 165:15 <b>telling</b> 78:17 84:6 122:5 <b>tend</b> 77:12 <b>tends</b> 10:13 60:21 99:25 <b>term</b> 13:7 52:7 59:18 73:4 129:18 140:15 142:19 156:25 <b>terms</b> 10:5 38:11 39:8 59:21,21 93:3 119:25 144:22 <b>terrifying</b> 84:4 129:7 <b>test</b> 159:6 <b>testimony</b> 19:13 96:14 97:20 113:21 167:2,3 <b>texas</b> 3:19 <b>text</b> 155:19 <b>texts</b> 16:20 <b>thank</b> 9:8 17:25 148:9 169:16,19 170:24 <b>theoretically</b> 21:23 49:4 140:8 155:2 <b>thereof</b> 173:12	<b>thing</b> 6:22 13:24 23:5 39:23 63:8 63:10,12 67:16 75:17 80:20,21 102:15 112:6 120:16 168:19,21 171:8 <b>things</b> 6:18 12:24 13:22 18:11,19 19:20 23:7,22 25:4 26:6,10,12 30:9 33:7,9 36:5 38:15 41:2 44:22 47:10 49:20 59:24 60:7,15,19 72:6 78:11 79:11 80:8 90:11 94:10 98:13 99:6 106:4 109:16 123:2,9 126:17 139:16 148:5 150:24 151:10 157:25 158:1,8 161:3 165:10 167:16 168:3 169:10 171:11 <b>think</b> 5:17 7:1,12 14:5,6,13 16:13,14 20:2,8,19,21 24:20 27:25 31:3 32:20 32:22 33:8,10 38:20 43:18 44:15 50:9,11 52:22 55:18 56:2,4 57:23 59:19 60:12 62:19 63:15 64:2 64:20,23 65:15,20 65:21 66:19,22 67:11,16 68:13 71:6 75:3 77:11 78:17 84:14 86:6 93:3 97:2 100:20	101:16,25 102:25 104:14 105:4 113:10 114:15,16 117:22 122:13 125:11 127:2,18 135:5 136:11 142:8,10 143:23 144:9,16 146:20 148:8 150:15,23 151:17 152:1 160:6,17 161:2 162:11 163:6 168:24 169:7,8,15 169:21,23 170:1 170:19 171:9,11 <b>thinking</b> 34:9 41:7 55:20 64:22 88:8 130:21 152:16 <b>third</b> 3:7 5:23 6:5 9:23 10:20 11:22 13:2,5,25 14:9 17:17 18:4,16 19:7,25 46:6 62:11,13 76:1 79:5 81:17,20 106:5,7 108:25 112:23 119:21 123:17,20,23 124:7,19 125:6 139:24 <b>thought</b> 19:20 96:16 166:9 167:12 <b>thousand</b> 86:22 <b>thousands</b> 8:23 38:4 86:11 97:10 <b>threads</b> 121:6,7 <b>three</b> 49:20 56:20 62:25 124:24 149:4 163:20
---	--	---	--



## HIGHLY CONFIDENTIAL

[throw - underline]

Page 34

<b>throw</b> 121:21 <b>thursday</b> 1:14 2:23 4:2 <b>tie</b> 42:3 <b>tied</b> 23:7 <b>tier</b> 169:12 <b>time</b> 9:8 10:20 16:14 22:7,10 24:24 31:8 33:11 42:15 50:2,8 60:14 61:17 67:18 69:23 70:25 71:24 72:10 85:2 95:20 101:3 107:17 119:17 125:17 135:14 149:2 151:7 152:9 153:23 163:10 164:7 173:7 <b>timeline</b> 17:2 54:12 66:22 116:19 124:17 131:14 135:16,19 151:11,12,15 <b>timelinedb</b> 54:6 158:15 <b>timely</b> 58:23 <b>times</b> 28:8 29:10 32:8 95:18 158:7 <b>timestamp</b> 39:19 <b>titan</b> 110:12 121:10 131:9 <b>tls</b> 108:5 <b>today</b> 4:7,10,16,20 5:4 6:22 7:5,10 8:1 15:20,23 30:20 32:5,18 34:17 56:25 57:1 57:4 80:6 112:21 130:6 136:18 140:9 159:21	168:1 171:9 <b>today's</b> 170:25 172:10 <b>token</b> 10:14,21 11:2,3,6,9,12,19 12:7 19:2 39:6,9 79:15 80:3,4,9,18 80:19,22,24 81:4,8 81:12,12,16 117:9 <b>tokens</b> 11:17,17 80:7 109:15 137:21 <b>told</b> 30:9 88:21 167:25 <b>ton</b> 23:23 71:12 <b>tool</b> 56:16 97:21 114:6 117:13,19 121:7 138:5,23 157:23 161:15 <b>tools</b> 36:4,7,9 130:5 <b>top</b> 38:4,9 67:1 73:19 94:25 107:12 110:8 114:22 127:3 147:12,20,24,25 149:3 153:21 155:6 <b>topic</b> 8:10 57:11 61:12 105:4 <b>topics</b> 8:13,15 54:21 <b>tos</b> 111:3,17 <b>tossed</b> 135:14 <b>total</b> 120:8 <b>touch</b> 97:13 142:11 <b>touched</b> 26:24 31:8 <b>trace</b> 56:24 61:21 62:17 89:7,9	102:4 103:3 141:15 <b>tracing</b> 78:1 <b>track</b> 18:2 26:22 28:9 29:21 36:5,9 58:15 77:7 85:7 93:23 101:20 <b>tracked</b> 63:20 71:20 75:23 <b>tracking</b> 36:25 69:21 70:4,21 78:22 <b>traffic</b> 33:16 110:25 <b>traits</b> 78:5 90:24 93:17 <b>transaction</b> 67:18 68:17 <b>transactions</b> 67:18 68:13,19 <b>transcribed</b> 173:8 <b>transcript</b> 1:12 68:4 169:22 <b>transient</b> 99:4 <b>transition</b> 156:21 <b>transitioned</b> 156:19 <b>traverse</b> 105:16 <b>traverses</b> 61:24 <b>trick</b> 101:15 <b>tried</b> 43:25 63:17 <b>triggered</b> 142:15 <b>true</b> 11:15 87:6 90:6 128:4 141:6 <b>truth</b> 7:9 <b>try</b> 26:10 58:8 73:5 169:10 <b>trying</b> 13:24 18:2 27:16 29:3 30:1,3 30:13 32:11 33:12 34:11 45:9 46:22	50:20 56:7 74:2 89:7,9 92:7 101:14 103:3 111:5 116:9 119:11 120:7 141:14 <b>tupperware</b> 100:18 <b>turn</b> 78:2 <b>turned</b> 70:22 <b>tw</b> 100:14,17 158:16 <b>two</b> 6:18 7:19 10:25 12:2 30:2 62:25 80:7 81:14 105:16 111:18 121:2 123:2 135:15 138:9,13 166:1 <b>type</b> 11:17 41:5 46:1 71:2,7,8,10 80:9 97:6 108:23 109:15 137:4 <b>types</b> 9:13,18 16:6 25:3 44:21 74:23 76:10,15 <b>typewriting</b> 173:9 <b>typically</b> 87:22 94:5 149:4
<b>u</b>			
<b>udb</b> 93:2,3 96:3,8 96:25 97:4 158:19 <b>udbs</b> 158:25 159:2 <b>uh</b> 68:14 <b>ui</b> 108:11,12,14 109:1 130:11 <b>uid</b> 77:2 <b>un</b> 8:18 <b>underline</b> 4:15 172:4			

## HIGHLY CONFIDENTIAL

[underlying - user]

Page 35

<b>underlying</b> 7:8 71:22 77:16 87:12 95:20 <b>underneath</b> 38:3 117:4,11 162:12 <b>underscore</b> 106:20 106:21 107:10,10 109:22,23,23,24 109:24,24 110:7 110:13 <b>understand</b> 8:24 10:7 13:23,25 15:21 20:19 27:16 29:3 31:14 45:11 46:22 71:13 78:2 80:25 84:5 86:25 90:4 111:6 116:24 119:8 120:8 143:1 143:11 145:8 155:2,8 168:23 <b>understanding</b> 9:2 28:1 50:24 82:13 83:3 85:2 105:4 106:10 117:22 119:5 121:19 123:9 138:17 159:12 161:1 <b>understood</b> 134:14 136:19 139:9 143:8 172:6 <b>undertaking</b> 138:13 <b>unfortunate</b> 78:25 <b>unfortunately</b> 42:16 62:9 100:12 120:14 <b>unicorn</b> 44:13,14 44:15,16,19 45:1,3 99:8 155:21 166:13	<b>unicorns</b> 167:24 168:1 <b>unifying</b> 68:24 <b>unintentionally</b> 41:22 <b>unique</b> 30:11 60:5 61:6 125:20 162:8 <b>unit</b> 29:1 30:4 48:9 48:22 49:25 77:3 126:23 <b>units</b> 28:12 <b>universe</b> 22:24 <b>unrelated</b> 13:18 150:3 <b>unsecure</b> 11:18 <b>up2x</b> 160:8 <b>update</b> 152:9 169:2 <b>updates</b> 74:22 <b>upload</b> 76:1,3,3,13 76:16 <b>uploaded</b> 94:15,21 <b>uploading</b> 74:9 75:14,16 84:16 116:17 <b>url</b> 74:24 <b>usage</b> 129:24 <b>use</b> 5:1 11:16 19:3 26:18,20,21 28:22 33:25 38:1 40:18 42:11 51:19 54:7 55:3 56:25 57:4 58:18 62:7,20 63:3,5,11 64:5 65:12 67:8,10 68:21 70:7 71:2 72:25 75:14 76:5 76:20 77:14,18,23 77:24 78:1 79:11 79:19 81:15,19 82:16,22 86:1	87:20,21,22,25 91:18,22 92:24 93:20 95:21 96:10 96:12 99:23 102:9 103:8 107:24,24 109:13 116:10 121:20 123:15 126:6 129:12,20 130:6,9,16,17 133:18 134:25 140:7 152:21,21 153:3,7 157:21,25 161:2 165:19,21 165:23 <b>useful</b> 34:3 41:25 105:5 <b>user</b> 1:3 2:3 4:15 4:21,22,24 5:2,13 7:14 9:14 10:1,20 10:23 11:2,17 12:14 13:11 14:1 14:1,10,10 15:1,22 17:18 18:4,7,11,24 19:2 20:25 21:3 22:1,4 23:9,10,21 24:15,15,17 25:20 25:21,22 26:8,13 26:17 27:17 28:24 29:6,9,10,12 30:4 30:5,14 31:5,8,10 31:11,16 32:5,18 34:4 35:11,25,25 36:6,9,13,14,17,22 37:1,14,14 38:17 38:23,24 39:4,6,11 39:15,20 40:12,13 40:15 41:19 42:2 42:13 43:7 45:1 45:15,18 47:16,18 47:20 48:4,6,13 49:4,15 50:22	51:2,11,17,17,19 51:22 52:1,5,6,7 52:10,11,12,14,15 52:23,24,25 53:5,5 53:11,11,12,13,19 53:19,25 54:6,9,15 58:18 59:19,21 60:5,8,10,11,15,16 60:16,23 61:1,2,4 61:5,7,20,23,23 63:12,18,20,24,25 64:15,23 65:2,10 65:16,18,18,21,22 65:25 66:4,8,9,11 66:24 67:9,11 68:24 69:2,3,3,20 70:4,16,17,22 71:16,19 72:16 73:13 74:1 75:8 75:15 76:13 77:2 78:20 79:2,2,5,12 79:15 80:3,9,11,15 80:17,19,20,22,23 80:25 81:11,13,16 81:21,25,25 88:23 89:2,21 90:15,20 91:3,13,24 93:24 94:5,13,15,17,19 95:13,13,25 96:24 98:7,11 99:15 103:22 105:10,12 105:14,15,17,20 105:22 108:19 110:2,4,4,17 112:14,17 114:12 118:21 119:13,14 120:2,10 121:5,6 121:15 122:14,21 122:24,24,24 123:1,2 124:14 125:20 130:7,13
--	--	--	--

## HIGHLY CONFIDENTIAL

[user - websites]

Page 36

131:13,24 134:13 134:20 135:6,14 141:18 142:18,19 143:1,14,21 144:5 144:15 146:13,20 146:20,24,25 147:3,21 150:4,11 150:12,14,19,25 151:25 152:16 153:25 155:2 158:4,20,21,23 159:1 161:9,14 164:7 165:11 166:2 <b>user's</b> 5:18,21,24 6:3,7 11:23 12:3,8 12:18 15:17 45:24 46:5,7,8 48:8 53:9 62:9,12,13 69:11 74:3 82:8 90:18 90:22 102:3 118:18 119:20,22 119:23 123:13,22 124:1,12,15,20,24 143:6 152:17 <b>users</b> 5:3 19:3 29:3 32:9 33:2 35:19 39:3,7,13 45:7 49:16 53:9 59:15 60:10 64:14 69:20 74:7 75:4 76:20 77:21 78:4,16 88:3,10 90:15 103:20 106:24 107:3,7 110:7 139:14 150:7,9 151:23 161:5 <b>uses</b> 17:12 56:8 57:10 78:2 79:15 94:23 123:11 158:2 162:12	<b>usually</b> 15:6 33:15 62:22 64:25 108:11 <b>v</b> <b>value</b> 170:19 <b>values</b> 58:17 <b>variety</b> 24:20 37:10 <b>various</b> 60:22 63:1 158:8 <b>vary</b> 99:24 <b>varying</b> 128:7 <b>veer</b> 142:7 <b>veered</b> 159:25 <b>verified</b> 81:7 <b>verify</b> 20:24 119:4 <b>version</b> 12:16 26:2 26:3 150:20,22 151:8 <b>versioning</b> 149:5 <b>versions</b> 149:1,5,9 <b>versus</b> 9:15,21,24 24:22 30:10 33:23 47:18 132:8 134:12 153:5 <b>verticals</b> 57:6 <b>video</b> 58:3,5 94:10 152:9 <b>videos</b> 63:7 126:17 135:12 156:1,18 <b>view</b> 29:7 36:13 106:24 114:7 <b>views</b> 20:25 21:6 <b>violate</b> 18:20 <b>violates</b> 19:9 <b>violations</b> 37:10 <b>virtually</b> 2:24 <b>visible</b> 107:3 <b>visits</b> 71:19 <b>vocabulary</b> 37:22	<b>volume</b> 99:25 128:14 <b>volumes</b> 100:3 <b>voluminously</b> 119:6 <b>volunteer</b> 44:10 <b>w</b> <b>wait</b> 21:19 55:11 <b>walk</b> 14:4 <b>walking</b> 14:7 <b>wall</b> 114:3 <b>want</b> 5:12 7:25 13:22 20:24 21:16 21:18 23:5 35:2 37:23 40:21 46:16 50:6 52:11 53:1 55:8,9 56:1 67:18 73:19 74:23 89:24 98:19 100:21 104:10,16 105:10 106:14 119:17 124:10 133:23 145:20 162:19 164:20 166:21,23 170:14 172:7 <b>wanted</b> 5:8 91:17 91:19 93:23 94:2 120:12 172:6 <b>wants</b> 102:22 142:9 <b>warm</b> 160:11 <b>washington</b> 3:7 <b>waste</b> 50:8 119:17 <b>watch</b> 135:12 <b>watched</b> 135:6,14 135:20 <b>watching</b> 125:5 139:18 <b>water</b> 55:3 <b>way</b> 9:17 10:10 15:3,18 18:21,23	19:11 24:20 35:11 35:19 43:15 45:11 45:14 49:21 51:4 52:22 66:11,18 69:2 72:25 75:3 75:23 89:20 92:7 95:3 96:23 98:13 98:20 100:18 105:7,20,24 115:25 121:24 122:1 131:25 133:11 134:19 135:4 136:3 147:21 158:13 165:25 173:12 <b>ways</b> 9:16 12:2 29:14 40:11 49:10 64:19 76:9 136:7 152:23 <b>we've</b> 23:15 24:3 24:25 34:25 54:21 56:2 110:19 117:22 134:19 142:18 152:1 <b>wealthy</b> 32:11 <b>weaver</b> 3:9 9:8 16:12,20 20:12 56:4 58:13 101:11 101:25 102:20 146:12 147:15,18 148:9,15 169:15 172:16 <b>web</b> 9:25 15:8 26:5,10 56:20 80:8 108:12 109:17 144:9 <b>webhooks</b> 74:21 75:1 129:3 <b>website</b> 73:9 76:3 <b>websites</b> 71:19
---	--	---	---

## HIGHLY CONFIDENTIAL

[welcome - zarashaw]

Page 37

<b>welcome</b> 155:17	118:24	14:8 16:4 17:25	12:1,6,12,17 13:8
<b>went</b> 63:17 115:2	<b>world</b> 17:8 61:10	18:9 21:11 22:21	14:6,20,25 15:12
163:9	85:18	24:6 28:16 31:13	15:19,24 17:9,19
<b>whereof</b> 173:13	<b>worried</b> 27:22	33:13 35:14 39:22	18:7,10,18,23 20:1
<b>whitelist</b> 97:17,19	<b>worth</b> 5:17 34:2	43:4,24 44:12	20:2 21:2,3,7,9,13
<b>whitelisting</b> 96:14	84:3 87:10 99:2	47:21 54:25 55:5	22:5,9,17,25 23:11
<b>whitelists</b> 16:6	142:2 153:13	55:7 60:4 66:3	24:19 25:9 26:24
97:21	<b>write</b> 16:23 32:14	68:7 73:22 83:12	27:13,18,25 28:13
<b>wide</b> 24:19 88:19	33:2 88:21 89:14	90:3 91:10 93:13	28:17 29:11 31:1
<b>wifi</b> 21:22 22:1	97:5 99:24 110:19	95:23,24 96:16	31:18 32:20 33:1
120:1	110:22 111:4	97:2,15 99:18	33:10,14,21 34:6
<b>willing</b> 34:13	115:18 116:9,11	104:6 105:1	34:23 36:8,15
102:7	116:24 125:5	109:22 115:14	37:7,16,25 38:8,13
<b>winding</b> 137:1	128:2	116:5 117:20	38:19 39:10,16
<b>windows</b> 152:8	<b>writes</b> 111:7	118:4,15 119:8	40:9,18 41:11,20
<b>wish</b> 46:4	115:16 116:4	121:23 122:6	42:14,24 44:2,9,13
<b>witness</b> 20:10	<b>writing</b> 40:24	124:6,9 126:8	44:16 45:16,19,25
173:13	72:21	128:12 129:4	46:10,24 47:8,11
<b>witnesses</b> 4:9 7:25	<b>written</b> 7:25 24:23	133:17 134:14	47:17,23 50:11
8:21 19:21	28:22 29:25 38:4	139:1 142:6 143:8	51:18 52:4 53:3,8
<b>wondering</b> 129:18	110:25 111:13,17	143:10 144:4,12	53:17 54:4 55:2
<b>word</b> 56:8 72:25	112:7 115:25	144:21 148:7	59:8,18 60:3,9
<b>words</b> 167:5	119:12 132:19	158:3 159:14,14	61:5,12,15,22 62:2
<b>work</b> 31:4 32:6	152:18	159:15 160:5	62:19 66:21,25
36:16 39:18 40:14	<b>wrong</b> 82:20 108:8	162:21 163:22	67:6,13 68:11
50:19 59:25 63:3	112:2,11 136:13	164:4 169:5	69:1,10,12,22
66:23 72:6 75:18	167:7	<b>year</b> 86:10 98:2	70:12,13,19 71:5
85:25 97:12	<b>wrote</b> 18:1 84:11	149:5	71:21 73:3,15
105:24 120:14,18	<b>x</b>	<b>years</b> 26:9 34:20	74:8,13 75:10,17
120:22 126:3,12	<b>x</b> 107:14 148:1	109:11 118:24	75:25 76:9 77:6
<b>worked</b> 23:24	<b>xdb</b> 93:2,4 96:5,6	137:17 139:2	78:7,21 79:7 81:4
28:13,14 32:10	96:8,19,21,25 97:4	<b>yesterday</b> 79:23	81:10,19 82:11,15
67:17 70:11 82:12	98:1 160:20	96:13 97:19	83:9,12,13,21 84:3
104:17 139:25	<b>xmpp</b> 107:10,14	<b>yesterday's</b>	84:14 85:5 87:4
140:23	107:16	113:20	87:10 88:14 89:3
<b>working</b> 8:18	<b>xpp</b> 112:1	<b>yield</b> 127:19	89:14 90:8 91:6
36:18,21 72:10	<b>y</b>	<b>yolo</b> 173:2	91:16 92:1,3,16
135:5	<b>yards</b> 26:23	<b>z</b>	95:2 97:1,8 98:21
<b>workings</b> 171:23	<b>yeah</b> 8:2,20 10:3,4	<b>zarashaw</b> 6:17	99:2,22 100:6,12
<b>works</b> 23:24 50:12	10:18 12:5,11	7:21 8:9 9:11,20	100:15 102:11
50:15 104:1,8		10:9,17,19 11:16	104:3 106:13

HIGHLY CONFIDENTIAL

[zarashaw - zoom]

Page 38

107:13,25 108:6	<b>zarashaw's</b> 8:12
108:22 109:5	138:18
110:10 111:18,22	<b>zealous</b> 7:9
112:4,22 113:3	<b>zippydb</b> 160:21
115:2,9,20 116:3,8	161:1
117:3,17 118:10	<b>zone</b> 22:8
118:11,23 120:13	<b>zoom</b> 1:11 3:3,14
121:8,19,25 122:9	59:3 80:14 155:17
127:25 128:20,24	
129:3,6,23 130:1,8	
130:19,24 131:4	
131:18 132:2,6,10	
132:17,20 133:1,5	
133:10,15,22	
134:3 135:4 136:7	
136:11 137:12,18	
138:6,12,22 139:7	
139:22 140:10,19	
141:1,11 143:24	
144:1,16 145:12	
145:13 146:16,22	
146:23 147:14	
148:3,10,21,24	
151:1,3,21 152:3,6	
152:13,19 153:2,9	
153:15,18,23	
154:5,11,15,23	
155:4,10,19,25	
156:5,8,12,17,23	
157:7,11,14,17,20	
157:24 158:6,11	
158:17,20,24	
159:19 160:9,12	
160:23 161:16,25	
162:9,23 163:3,7	
163:15,19,25	
164:7,15,24 165:5	
165:14,24 166:11	
166:16 167:8	
168:3 169:1	
170:24	

# Exhibit M

GIBSON DUNN

Gibson, Dunn & Crutcher LLP  
555 Mission Street  
San Francisco, CA 94105-0921  
Tel 415.393.8200  
www.gibsondunn.com

Rosemarie T. Ring  
Direct: +1 415.393.8247  
Fax: +1 415.801.7358  
RRing@gibsondunn.com

HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

March 7, 2022

VIA JAMS ACCESS

Special Master Daniel B. Garrie  
DGarrie@jamsadr.com

Re: *In re Facebook Consumer Privacy User Profile Litigation*, JAMS Ref No.  
1200058674

Dear Special Master Garrie,

Pursuant to the Hearing Order Regarding Plaintiffs' Motion To Compel Production of Plaintiff Data dated February 21, 2022, we submit this response on behalf of Facebook. We also thank the Special Master for granting an extension of our time to respond to allow us to provide responses that we believe would assist the Special Master in expediting resolution of the issue set forth by Judge Corley in her January 12, 2022 order: "[W]hat, if any, data from [systems other than DYI] should be produced consistent with Federal Rule of Civil Procedure 26(b)."

**I. USER DATA IN DYI**

The Special Master's first two questions seek information about what user data is included in DYI. We welcome the opportunity to address that issue to correct and clarify statements made throughout these proceedings which seem to have created the impression that the DYI system is limited to the first category of "discoverable user data" identified by Judge Corley in Discovery Order No. 9. As explained below, the DYI system includes *all three categories* of data: (1) data collected from a user's on-platform activity, (2) data obtained from third parties regarding a user's off-platform activities, and (3) data inferred from a user's on or off-platform activity. So the issue is whether Facebook should produce *more data* in categories (1), (2), and (3) from systems other than DYI, which we understand is the focus of the Special Master's remaining questions.



## GIBSON DUNN

Special Master Daniel B. Garrie  
March 7, 2022  
Page 2

HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

**Question #1.** Engineer documentation for the process of generating the DYI file for a Facebook user, including data flow diagrams that explain how the data is retrieved, data schemas, and the individual fields for each data type in the schema.

The following DYI engineer documentation is attached as **Exhibit A**.

- **Your Information Schema (DYI).** This schema describes each of the categories and types of data included in DYI.
- **Eng Guide: Adding to DYI Wiki.** This document provides step-by-step instructions for engineers from product teams to add data to DYI, including instructions regarding where data should live in the DYI taxonomy and DYI schema fields.

**Question #2.** A statement clarifying whether information available via the DYI tool “includes user data provided by third parties to Facebook, e.g., data relating to users’ off platform activity.”

Yes, DYI includes user data provided by third parties to Facebook, which is the second category of “discoverable user data” identified by Judge Corley in Discovery Order No. 9. Since this data is stored as raw logs of event data in Facebook’s data warehouse, Facebook had to build data pipelines from the data warehouse to production to (a) aggregate and group this event level data on a per user basis sorted in chronological order and (b) transform raw logs of event data in the data warehouse into a human readable information. This effort required complex and time-extensive engineering efforts.

This question from the Special Master highlights a broader factual issue, which we believe should be clarified to assist the Special Master in resolving the issue before him. In Discovery Order No. 9, Judge Corley identified the following categories of “discoverable user data”: (1) data collected from a user’s on-platform activity, (2) data obtained from third parties regarding a user’s off-platform activities, and (3) data inferred from a user’s on or off-platform activity. Plaintiffs repeatedly have asserted that DYI contains *only* data in category (1), citing a statement by Facebook’s counsel during a status conference before Judge Corley.<sup>1</sup> In fact, the full exchange with Judge Corley makes clear that Facebook’s counsel was explaining that the DYI tool does not contain *all* data in categories (1), (2) and (3), which is further demonstrated by the answer to this question confirming that DYI includes

---

<sup>1</sup> Pls’ Sept. 28, 2020 Mot. Compel at 7, Dkt. 526; Pls’ Oct. 18, 2021 Mot. Compel Production of Named Pls’ Content And Information at 3.

*(Cont’d on next page)*



## GIBSON DUNN

Special Master Daniel B. Garrie  
March 7, 2022  
Page 3

## HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

data in category (2). Plaintiffs also have cited documents produced by Facebook that they claim prove Facebook's production is limited to *only* data in category (1), most notably, a document referring to "Native Data," "Appended Data," and "Behavioral Data."<sup>2</sup> Again, this is incorrect. **Exhibit B** attaches the primary document Plaintiffs cite for this assertion and a table that describes how the Native, Appended, and Behavioral Data described in that document is either included in DYI or relates to features that were deprecated before this case was filed.

For completeness, we explain below the data in categories (1), (2), and (3) that are included in DYI, and attach as **Exhibit C** summaries of the data that has been produced for each of the named plaintiffs.

**Category (1): Data collected from on-platform activity.** DYI includes data provided by users and data observed by Facebook on the platform.

*User-provided data* includes profile data, user-generated content (e.g., posts, videos, photos, comments, stories), message content, friends, location check-ins, linked accounts in the Facebook family of products, and language choices.

*Observed data* includes clicks, profiles, Pages, Groups, and Events a user has visited, usage data, device data, networks and connections, data about user's activity level, advertisers with which the user has interacted, pages (user pages, pages a user liked or recommended, pages a user follows, pages a user has unfollowed), IP address when sending a message, users that a user has chosen to "see less" or "see first" in News Feed, time spent watching from a page, people whose profile a user has visited, last location, last active time, whether a user viewed someone's birthday story, people a user blocked on Messenger, page notifications, pages a user recommended, time zone, email address verification, Marketplace notifications, and interactions.

**Category (2): Data collected from off-platform activity.** DYI includes information provided to Facebook by third-party advertisers, app developers, and publishers about user interactions.

User interactions are things like opening a third-party developer app that integrates Facebook business tools, and visiting websites that integrate the Facebook business tools providing information about the user viewing content, searching for items, adding an item to a shopping cart, or making a purchase. Third parties share this off-platform activity with Facebook using "business tools," which are technologies designed to help website owners

---

<sup>2</sup> Pls' Sept. 28, 2020 Mot. Compel at 8, Dkt. 526.

# GIBSON DUNN

Special Master Daniel B. Garrie

March 7, 2022

Page 4

## HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

and publishers, app developers, advertisers, and others, integrate with Facebook, understand and measure their products and services, and better reach and serve people who use or might be interested in their products and services. Business tools include Facebook Pixel, Facebook SDK, Conversions API, Offline Conversions, and the App Events API. Facebook also receives impression data through Facebook Social Plugins (e.g., Like and Share buttons) and Facebook Login.

Facebook also receives contact lists with user information uploaded by advertisers for the purpose of serving ads to those users through what are referred to as custom audiences. These lists include hashed information about users (e.g., email addresses, phone numbers) that is used to try to match people on the list with Facebook user accounts; matched users are placed in the custom audience. This hashed information is discarded after the matching process.

**Category (3): Data inferred from a user’s on or off-platform activity.** DYI includes data that is derived from a user’s on or off-platform activity.

DYI includes information regarding ads interests; music recommendations based on genres of music a user has interacted with on Facebook; “your topics,” which is a collection of topics determined by a user’s activity on Facebook that is used to create recommendations for users in different areas of Facebook such as News Feed, News, and Watch; primary location; primary public location; friend peer group; creator badges (including labels like “visual storyteller” or “conversation starter” based on activity in Groups); time zone; language preferences (including preferred language for videos, languages you may know, preferred language); and mobile service provider and country code.

**Plaintiffs’ DYI Data.** Summaries of the DYI files for each named plaintiff showing that they include all three categories of “discoverable user data” identified in Discovery Order No. 9 are attached as **Exhibit C**.

## II. OTHER USER DATA AND SYSTEMS

The Special Master’s remaining questions focus on understanding what other systems contain user data and, as Judge Corley stated the issue, “what, if any, additional data should be produced consistent with Federal Rule of Civil Procedure 26(b).” To be sure, the DYI file does not include all data related to users, but that does not mean that production of that data is consistent with Rule 26. For example, as explained above, DYI includes data received from third parties regarding a user’s off-platform activity on apps and websites, such as viewing content and adding an item to a shopping cart, but does not include data identifying

## GIBSON DUNN

Special Master Daniel B. Garrie  
March 7, 2022  
Page 5

### HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

the specific content that was viewed or the item that was added to a cart for reasons that engineers will be prepared to explain at the hearing.

**Question #3:** A statement identifying systems that coordinate and schedule jobs that run against the Hive (a process that accesses table data in the Hive and aggregates user data to produce a meaningful data set). For each job that may involve user data, Facebook is to describe the data extracted in the job and where the job saves the data.

The majority of batch data processing of Hive data at Meta is handled by a system called Dataswarm, which is described below. The remaining minority of batch data processing is coordinated by FBLearder, which is a similar system derived from Dataswarm.

Dataswarm works by having employees (1) define atoms of computation called tasks and then having employees (2) explicitly state the dependency relationships between these tasks so that the system can initiate a task's computations after the preceding tasks have completed their execution. These tasks are treated as black boxes: the system knows nothing about what the task does beyond the rough type of computation performed. For any given task, Dataswarm does not know what data is used as inputs to the computations it orchestrates or what data is produced as outputs by these computations.

Facebook's current approach for identifying what data is consumed as inputs by a job and is generated as outputs by that job is a time-consuming manual process. Because Dataswarm performs millions of tasks each day, it is not possible to complete this manual process for all Dataswarm tasks. To respond to the Special Master's request, Facebook completed this manual process for a sample of 10 tasks run in Dataswarm on February 15, 2022. This sample is attached **Exhibit D**. Approximately five million Dataswarm tasks were run on February 15.

**Question #4:** A statement identifying the internal identifiers Facebook uses to track users across the Facebook platform, including a description of which Facebook systems use each identifier, how each system uses each identifier, and how Facebook maps identifiers to users.

Facebook primarily uses four types of internal identifiers for user data: (1) a user identifier (UserID), (2) Replacement ID, (3) Separable ID, and (4) App Scoped Identifiers. Each is described in more detail below.

**UserID.** Facebook uses an industry-wide technique called pseudonymization to represent users on the Facebook platform. In essence Facebook creates a canonical unique

# GIBSON DUNN

Special Master Daniel B. Garrie  
March 7, 2022  
Page 6

## HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

identifier that encapsulates information about the user (such as First Name, Last Name, email, phone numbers, etc). The encapsulation can be accessed by an identifier called a user identifier (or UserID); this is similar to a row being stored in a database table with the primary key being the userid and information of the user being values in the other columns. UserIDs are unique in Facebook's systems, such that two users cannot have the same UserID, and they are not recycled, meaning even after a user deletes their account, no other user can have the same UserID. The User ID is the canonical identifier to represent a Facebook user and is used in nearly all Facebook systems.

**Replacement ID (RID).** The RID is an identifier that supports Facebook's deletion practices by irreversibly disassociating data from a user. Every user is assigned an RID for the lifetime of their account. In data systems that do not support deletion (e.g. Hive), any user data retained for more than 90 days can only be retained with an RID. When a user deletes her account, Facebook deletes the record connecting the UserID to the RID so that data stored with that RID can no longer be connected to that user. Like the UserID, the RID represents a single user. Two users cannot have the same RID, and RIDs are not recycled. RIDs are used in Hive.

**Separable ID (SID).** The SID is similar to the RID, but allows Facebook to permanently disassociate Off Facebook Activity data from a user. Data Facebook receives from third parties about a user is associated with an SID (rather than UserID), and Facebook maintains a separate mapping between SIDs and UserIDs that can be accessed when data is processed. Through Facebook's Off Facebook Activity tool, users are able to clear their Off Facebook Activity. When a user does this, Facebook removes the mapping between the users' SID and UserID, which irreversibly dissociates the data stored with an SID from the user. Facebook then generates a new SID to be associated with the user's account moving forward. SIDs are used in Hive. More information about SID's is available at this link: <https://engineering.fb.com/2019/08/20/data-infrastructure/off-facebook-activity/>

**App-Scoped ID (ASID).** The ASID is an identifier that is sent to the third-party developer when a Facebook user has chosen to use the Facebook Login product to login to the services of the third-party developer. The App-Scoped ID also serves as a privacy protective identifier that Facebook created with the purpose of preventing different third-party apps from sharing data or amassing profiles of users across apps. This is why the App Scoped ID is created on a per-user per-app basis, such that different third-party developers do not receive the same identifier for the same Facebook user. ASIDs are stored in TAO.

## GIBSON DUNN

Special Master Daniel B. Garrie  
 March 7, 2022  
 Page 7

HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

**Question #5:** A statement identifying the sources from which Graph API pulls user data, including a high-level description of each source and the engineers that own each source.

**TAO.** The Facebook production infrastructure is centralized around a distributed data store for the social graph, called TAO (The Associations and Objects), which is the primary source from which Graph API pulls data (including user data). TAO is a high-performance service for storing, caching, and querying the graph for nodes and associations, by providing a clean interface for internal and external developers to integrate into the social graph, abstracting away many of the complexities of developing and maintaining a data storage at scale.

- **For more information about TAO:** <https://engineering.fb.com/2013/06/25/core-data/tao-the-power-of-the-graph/>
- **Technical Point of Contact:** [REDACTED]

TAO provides a layer of abstraction to describe objects and relations through EntSchema and Node. The underlying technology that is used to store these data models is a collection of technologies such as MySQL, Manifold, ZippyDB/Akkio, Memcache, and Laser. In limited cases, and for specific purposes, other systems may access these systems directly.

**MySQL:** MySQL is TAO's backbone. It provides transactional and availability properties to columnar data. For example, a user's comment can be stored in a MySQL database as a row in a table, where the comment id is the primary key and the comment is a text field. As another example, the fact that someone liked a comment can be represented by an association with the type like from the comment id and the user id, this could be represented as 3 columns in the table, with comment id, user id, and type of reaction.

- **Technical Point of Contact:** [REDACTED]

[REDACTED]

- **Technical Points of Contact:** [REDACTED]

## GIBSON DUNN

Special Master Daniel B. Garrie  
March 7, 2022  
Page 8

### HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

**ZippyDB/Akkio:** ZippyDB is a persistent key-value storage service that provides reliable, scalable, geo-replicated storage for high throughput applications. ZippyDB provides very efficient key value lookups, faster than MySQL but are not transactional in nature. Akkio operates on top of storage services like Zippy. Akkio splits the data sets into units with strong locality, which can then be geographically placed close to where they are most likely to be accessed.

- **For more information about ZippyDB/Akkio:**  
<https://engineering.fb.com/2021/08/06/core-data/zippydb/> (ZippyDB);  
<https://engineering.fb.com/2018/10/08/core-data/akkio/> (Akkio).
- **Technical Points of Contact:** [REDACTED] (ZippyDB), [REDACTED]  
[REDACTED] (Akkio)

**Memcache:** Memcache provides a cache to the TAO service, the cache ensures that queries that have recently run are available quickly through the cache, as the response time of a cache are orders of magnitude quicker than a lookup in a MySQL database.

- **Technical Point of Contact:** [REDACTED]

**Laser:** Laser is an indexing service that provides low-latency (typically a few milliseconds) read access to specific sets of Data Warehouse data. Data in the Data Warehouse is not indexed on a per user basis, but some Facebook products need to surface the calculations that the product performs to the production systems so a GraphAPI can request the data. Laser is only used to support first-party (i.e. Meta) products.

- **Technical Point of Contact:** [REDACTED]

There are other technologies that keep TAO operational that ensure the service is performing at a high-performance.

**Question #6:** Any agreements relating to user data that Facebook has with each of the following entities: Netflix, Microsoft, and YouTube.

Facebook will submit via email to the Special Master and Plaintiffs contracts relating to user data with Netflix, Microsoft, and YouTube that have been produced in this case. Facebook will separately submit to the Special Master contracts with these entities that have not been produced to date in camera, to provide Facebook an opportunity to fully analyze

GIBSON DUNN

Special Master Daniel B. Garrie  
March 7, 2022  
Page 9

HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

and comply with any notice provisions in these agreements. Once Facebook has completed this analysis and complied with any notice provisions, Facebook will produce the contracts to Plaintiffs, consistent with Facebook's agreement to search for additional contracts Plaintiffs recently requested.

Facebook notes that Meta Platforms, Inc., Netflix, Microsoft, and YouTube are some of the largest technology companies in the world. These entities are parties to a very large number of agreements, the vast majority of which are confidential business agreements and outside the scope of this litigation. Should Facebook identify additional agreements with these entities that are responsive to the Special Master's request, Facebook will provide them.

**Question #7:** Update list of 149 systems identified in the Declaration of David Pope.

In discussing the topics and questions to be addressed during the hearing on Wednesday, March 9, 2002, the engineers who will be attending the hearing have explained to counsel that the list of 149 systems and categorization discussed during the hearing with David Pope on January 14, 2002, are not an effective way of assisting the Special Master in understanding the user data that exists in Facebook systems. Instead, they approach this question from the perspective of a Facebook engineer building product experiences or conducting data analysis who think of user data in terms of production systems and the data warehouse. As these engineers will explain at the hearing, they look to TAO for production systems with narrow exceptions, and to Hive for the data warehouse.

Sincerely,

A handwritten signature in cursive script, reading "Rosemarie Ring".

Rosemarie T. Ring

# Exhibit N



HIGHLY CONFIDENTIAL - ATTORNEYS EYES ONLY

UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA

IN RE: ) MDL No. 2843  
)  
FACEBOOK, INC. CONSUMER ) Case No. 18-md-02843-  
PRIVACY USER PROFILE ) VC-JSC  
LITIGATION )  
----- )

HIGHLY CONFIDENTIAL

ATTORNEYS' EYES ONLY

SPECIAL MASTER: DANIEL GARRIE

REMOTE HEARING  
(Via Zoom Videoconference)  
Wednesday, March 9, 2022

REPORTED BY: Michelle Milan Fulmer  
CSR No. 6942, RPR, CRR, CRC

Page 1

UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA

IN RE: ) MDL No. 2843  
)  
FACEBOOK, INC. CONSUMER ) Case No. 18-md-02843-  
PRIVACY USER PROFILE ) VC-JSC  
LITIGATION )  
----- )

Remote hearing taken before Michelle Milan  
Fulmer, a Certified Shorthand Reporter for the  
State of California, commencing at 10:20 a.m.,  
Pacific Standard Time, Wednesday, March 9, 2022.

1 APPEARANCES OF COUNSEL:

2  
3 FOR PLAINTIFFS:

4 BLEICHMAR FONTI & AULD

BY: Lesley Weaver, Esq.

5 Anne K. Davis, Esq.

555 12th Street, Suite 1600

6 Oakland, California 94607

TEL: (415) 445-4003

7 EMAIL: lweaver@bfalaw.com

adavis@bfalaw.com

8 (Via Zoom Videoconference)

9  
KELLER ROHRBACK, LLP

10 BY: Derek W. Loeser, Esq.

Cari Campen Laufenberg, Esq.

11 1201 Third Avenue, Suite 3200

Seattle, Washington 98101

12 TEL: (206) 623-1900

EMAIL: dloeser@kellerrohrback.com

13 claufenberg@kellerrohrback.com

(Via Zoom Videoconference)

14  
15 FOR DEFENDANT:

16 GIBSON DUNN & CRUTCHER

BY: Martie Kutscher Clark, Esq.

17 1881 Page Mill Road

Palo Alto, California 94304-1211

18 TEL: (650) 849-5348

EMAIL: mkutscherclark@gibsondunn.com

19 (Via Zoom Videoconference)

20  
GIBSON DUNN & CRUTCHER

21 BY: Rosemarie Ring, Esq.

555 Mission Street, Suite 3000

22 San Francisco, California 94105

TEL: (415) 393-8247

23 EMAIL: rring@gibsondunn.com

(Via Zoom Videoconference)

1 FOR DEFENDANT:

2 FACEBOOK LEGAL

BY: Ian Chen, Esq.

3 Sandeep Solanki, Esq.

1 Hacker Way

4 Menlo, California 94025

EMAIL: ianchen0@fb.com

5 (Via Zoom Videoconference)

6 ALSO PRESENT:

7 Michael Mann

8 David Ko

(Via Zoom Videoconference)

1 Wednesday, March 9, 2022

2 10:20 a.m. - 2:28 p.m.

3 \*\*\*

4 MR. GARRIE: We will go on the record.  
5 We will take breaks for the court reporter's sake  
6 every 45 minutes to an hour basis. I lose track of  
7 time, Michelle. So you need to tell me because I  
8 get engrossed in something and time becomes  
9 irrelevant. So if you would just kindly remind me.  
10 It's not that I don't want to give anybody a break.  
11 It's just that I get -- I lose track of time. So I  
12 apologize.

13 So on the record, let's just test  
14 everybody's audio for the record.

15 Mr. Patel, just say a couple words.

16 MR. PATEL: Hello, everybody.

17 MR. GARRIE: Court Reporter, we're good?

18 THE COURT REPORTER: Yes.

19 MR. GARRIE: Okay. Mr. Zarashaw.

20 MR. ZARASHAW: Good morning.

21 MR. GARRIE: We're good.

22 Mr. Mitchell.

23 MR. MITCHELL: Good morning.

24 MR. GARRIE: Mr. Clark.

25 MR. CLARK: Good morning.

1 MR. GARRIE: All right. This is  
2 Special Master Daniel Garrie.

3 The purpose of today's hearing is to  
4 discuss, and only today, hopefully no further  
5 hearings are necessary, regarding named plaintiff  
6 data that was the order that was issued and  
7 attempting to identify and figure out what are the  
8 appropriate scope of data within the discoverable  
9 data pursuant to the court's order relating to the  
10 named plaintiff data that is to be produced in this  
11 litigation.

12 Today's format is I will largely be asking  
13 questions of the engineers. I have invited counsel  
14 for Facebook to object if I go into anything of a  
15 privileged nature as well as instructed the  
16 witnesses to ask if they have any questions, to ask  
17 any clarifying questions; if they think something  
18 they're saying may be privileged, to ask to speak to  
19 counsel. We have breakout rooms. We'll go  
20 accordingly and have those conversations.

21 With all of that said, I first want to just  
22 officially on the record recognize and thank  
23 Facebook for accommodating an accelerated timeline  
24 for a very complex set of problems. I don't want  
25 to -- I realize litigation has been going on for

1 some time. I realize there's a lot of history here,  
2 but I also recognize that two weeks isn't a lot of  
3 time to answer all of the questions, or even a month  
4 to answer in detail the technical questions that  
5 I've asked.

6 So the answers and data that has been  
7 provided is very helpful. I want to thank Facebook  
8 for making that effort, making the engineering  
9 resources available. It's certainly helping  
10 expedite things and move things along and I wanted  
11 to officially recognize that on the record.

12 With that said, counsel will not be asking  
13 questions. If there is a burning desire for counsel  
14 to ask a question, you can raise your hand and I  
15 will, when I'm done with my questions, invite  
16 questions accordingly and, if I think they're  
17 appropriate, then we will ask them. Nobody will ask  
18 the engineers questions, but me.

19 I did invite counsel for plaintiffs to  
20 provide questions and material and I reviewed them.  
21 Some of them I think are relevant. Some I do not.  
22 I will use them accordingly. I do want to thank  
23 plaintiffs as well for going out of their way to  
24 help identify appropriate questions as it relates to  
25 this issue we have before us now.

1 With all of that said, I'm going to jump  
2 right in to lay out the format of how I'm --

3 MS. WEAVER: I apologize for interrupting,  
4 but do you want to swear in the witnesses?

5 MR. GARRIE: I will. Thank you for  
6 reminding me. I often forget. There's usually a  
7 court reporter and I'm not the one usually doing it.  
8 So there's usually --

9 Would the court reporter swear in? I guess  
10 I don't know if we can group swear in, but whatever  
11 is comfortable, however the court reporter thinks  
12 best, swear in the witnesses.

13 THE COURT REPORTER: Okay. And,  
14 Mr. Garrie, are you going to be asking questions  
15 of Mr. Patel, and then the next time Mr. Clark,  
16 et cetera?

17 MR. GARRIE: Oh, no.

18 THE COURT REPORTER: Oh, okay.

19 MR. GARRIE: It's going to be a  
20 free-for-all. I'm not sure who knows what. So I  
21 don't want to -- they know way more about Facebook  
22 and the engineering systems than I ever will. So I  
23 don't know which one of the four will be best  
24 situated. So I'll direct the question to who I  
25 believe and they can go around and say, "I have



1 something to add."

2 THE COURT REPORTER: So, Mr. Zarashaw, will  
3 you raise your right hand to be sworn?

4

5 EUGENE ZARASHAW,  
6 called as a witness, having been first duly sworn by  
7 the Certified Shorthand Reporter, was examined and  
8 testified as follows:

9

10 THE COURT REPORTER: Mr. Patel, will you  
11 raise your right hand to be sworn?

12

13 MAYUR PATEL,  
14 called as a witness, having been first duly sworn by  
15 the Certified Shorthand Reporter, was examined and  
16 testified as follows:

17

18 THE COURT REPORTER: Mr. Mitchell, will you  
19 raise your right hand to be sworn?

20

21 BEN MITCHELL,  
22 called as a witness, having been first duly sworn by  
23 the Certified Shorthand Reporter, was examined and  
24 testified as follows:

25

1 THE COURT REPORTER: Mr. Clark, will you  
2 raise your right hand to be sworn?

3 I think he froze.

4 MR. GARRIE: When Mr. Mitchell rejoins, the  
5 JAMS -- Mr. Mitchell is here.

6 When Mr. Clark rejoins -- thank you,  
7 Mr. Mitchell -- the JAMS moderator hopefully is  
8 still here and will -- here he is. I'm admitting  
9 him now.

10 MS. RING: May I say, Special Master  
11 Garrie, I just want to be clear, I guess, in  
12 response to the free-for-all comment.

13 These witnesses have been prepared to  
14 address specific topics. So I just want everyone to  
15 be aware of that and I want to give them comfort  
16 that they're not going to be expected to guess to  
17 things that they don't know. Right.

18 MR. GARRIE: Well, I was going to -- to  
19 your point, Counsel Ring, what I was going to say is  
20 if any of the engineers -- it's completely okay not  
21 to know something. You guys have millions of lines  
22 of software code and areas of specialty and teams  
23 you lead and it's completely acceptable to say, "I  
24 don't know." I don't expect you to know everybody.

25 I just don't know if you worked on in the

1 past in certain areas or something else. But please  
2 say -- if you don't know something and it's  
3 completely unacceptable to answer, just say, "I  
4 don't know." I don't expect you to know the answers  
5 to everything and we'll go through it like that.  
6 Just if you have anything to add to a particular  
7 question, then, you know, you'll add; and if not,  
8 you won't.

9 Is that all right with you, Counsel Ring?

10 MS. RING: Yes. Thank you, Special Master  
11 Garrie. I just wanted to be -- you know, I want  
12 them to feel comfortable and they've been -- we've  
13 been talking to them for a long time and they're  
14 prepared to answer your questions. I just wanted to  
15 be clear about what we've prepared them for and  
16 discussed with them.

17 MR. GARRIE: Oh, 100 percent. Yeah.  
18 That's my expectation. I just -- they're the  
19 engineers and I'm not. So that's why I just sort  
20 of give them. And project mana- -- at some point  
21 we're engineers, maybe more or less than. I don't  
22 know.

23 But the point is, Mr. Clark, if we can just  
24 quickly swear you in, and then we can kick it all  
25 off.

1 MR. CLARK: We had a power outage. So I'm  
2 on my phone.

3 THE COURT REPORTER: Okay. Can you raise  
4 your right hand to be sworn, Mr. Clark?

5

6 MIKE CLARK,  
7 called as a witness, having been first duly sworn by  
8 the Certified Shorthand Reporter, was examined and  
9 testified as follows:

10

11 MR. GARRIE: Okay. We've gone through all  
12 the formalities.

13 With that said, before we jump into  
14 everything, the way I'm going to approach today is  
15 I've presented a series of questions. I'm going to  
16 go through the questions and I've read the responses  
17 and I have additional clarifying information I'm  
18 looking for regarding the responses to your  
19 questions. So if you have anything else you --

20 So what I'm going to do is I want to ask  
21 my, at least, one question, and then if you guys  
22 want to add any further clarifications or any other  
23 statements or other things, I'm going to welcome  
24 them. I do want to appreciate -- thank you for  
25 categorizing them.

1           And for the record, Counsel Ring, I also  
2   agree and recognize that a different way of  
3   approaching the data than looking at all 149 systems  
4   is definitely the way to go, but I wanted to finish  
5   off the initial task so we have a complete record  
6   for that, but I do fully agree with you and that's  
7   what I'm hoping we can start and finish today.

8           With that said, I'm going to read out just  
9   so we have a clear record of what I'm talking about  
10  so there's no ambiguity.

11           It says, "Accordingly, the court rules,"  
12  this is Judge Corley, "rules the discoverable user  
13  data at issue includes (1) data collected from a  
14  user's on-platform activity, (2) data obtained from  
15  third parties regarding a user's off-platform  
16  activities, and (3) data inferred from a user's on  
17  or off-platform activity."

18           So I just -- so when I say in the letters  
19  and the back and forth, there's one, two, and three  
20  constantly going back and forth, that's when we say  
21  one, two, or three what we're referring to so  
22  there's no ambiguity here.

23           With that said, I think we'll kick off with  
24  the first question and the first question really was  
25  talking about the DYI file and how it's generated

1 and the information, you know, and the answers, but  
2 what I saw -- sorry. I'm just looking at my  
3 outline. Actually I did quite a bit of homework  
4 myself. What I saw and in my review of it -- and I  
5 had a chance to read it. I luckily don't need a lot  
6 of sleep. But I had a chance to read it and, from  
7 what I could tell, it describes how data is tagged  
8 and to be included in the DYI export, but not how  
9 data is extracted by the DYI file generation  
10 process.

11 So, you know, what I'm trying to understand  
12 is where the data comes from and how it's filtered  
13 to only include tagged fields and what other fields  
14 it is to exclude.

15 So in Exhibit A, the access, in the  
16 engineering guide I think it's Page 29. Let me make  
17 sure we're talking apples to apples. Page 29.  
18 So on Page 29 in there, the engineering guide,  
19 there's this thing called Access your Information  
20 and can someone explain -- okay, I think I know  
21 who to ask, but I'll keep it broad for now --  
22 you know, explain how Access your Information is  
23 different from DYI?

24 Mr. Clark, we'll start with you.

25 MR. CLARK: Yeah. Of the transparency

1 products that we have, a DYI is allowing you to  
2 download your file. AYI is meant to be an  
3 interactive version that is live on the site and not  
4 downloading every single bit of information that's  
5 included in the DYI file. And so its function is  
6 slightly different in that it is set and populated  
7 with a similar, but not as a complete cut of data.

8 MR. GARRIE: DYI -- does AYI data--

9 I guess a way of thinking about it, is AYI  
10 small and then DYI contains all the AYI data?

11 MR. CLARK: Plus more.

12 MR. GARRIE: Well, yeah. So AYI and then,  
13 bigger, all of DYI has AYI plus whatever else.

14 MR. CLARK: That is correct.

15 MR. GARRIE: Okay. So then the AYI is  
16 just -- how is that -- is that working to do what?  
17 You said it was real-time data that they're looking  
18 at or --

19 MR. CLARK: Well, DYI is somewhat real  
20 time, but it takes time to cue up. I've got a long  
21 Facebook usage history, and so mine takes the better  
22 part of a couple days and is many gigs in size even  
23 with low-resolution photos.

24 AYI allows me to click through that  
25 information live on the site and be able to look at

1 that content on the Facebook site.

2 MR. GARRIE: It's basically re-skimming the  
3 data effectively.

4 MR. CLARK: Correct. And not putting it in  
5 a downloadable format.

6 MR. GARRIE: Yeah. Okay.

7 So then the reference is the Activity Log  
8 is a location for storing data.

9 I think, what is the Activity Log? I  
10 didn't see the definition of what an Activity Log is  
11 and I was hoping you could clarify what is the  
12 Activity Log.

13 MR. CLARK: Sure. There's a very -- in  
14 fact, they're very similar and some of the things  
15 that you click on in AYI actually is linked to a  
16 product called Activity Log. Activity Log is meant  
17 to be much more like the Facebook site itself as  
18 opposed to sitting outside of the site.

19 When you interact with the content in  
20 Activity Log, you can in some cases delete it or  
21 archive it right from the Activity Log. It is a  
22 subset of the data that's in AYI, but it is meant to  
23 be the logging of your interactions and the content  
24 that you've uploaded and those specific things  
25 inside of Activity Log that you can interact with



1 more so than you can in AYI. AYI is meant to read  
2 only. Activity Log I can actually interact with the  
3 content.

4 MR. GARRIE: Okay. I'm just writing it  
5 down.

6 So then where is the Activity Log data  
7 itself stored?

8 MR. CLARK: In production storage systems.

9 MR. GARRIE: So I think you gave me every  
10 single system you have.

11 When you say production systems, you're  
12 referring to a particular one? I know it's  
13 aggregated, backed up, and distributed.

14 So when you say production systems, can we  
15 define what is a produc- -- you have -- production  
16 systems becomes this very large list of things which  
17 I don't really need to touch on as long as it's the  
18 same and there's nothing new in the production  
19 system, if that makes sense.

20 MR. CLARK: Yeah. I have a response up on  
21 my screen; but since I can't see it, I'm going to  
22 defer to Mayur to help answer and clarify, Mr. Patel  
23 to help answer and clarify your question on, like,  
24 where some of that content may be.

25 MR. PATEL: All right. So just to answer

1 the question, to understand the question, you're  
2 asking where does Activity Log data reside, in which  
3 specific data store?

4 MR. GARRIE: Yeah. Like where is it  
5 actually stored?

6 MR. PATEL: Yeah. So it's stored in -- if  
7 you have the 149, it's stored in [REDACTED] In  
8 order to make it effective to produce the results in  
9 real time, we need to [REDACTED]

[REDACTED], and so we have a specialized store  
11 called [REDACTED]

12 MR. GARRIE: Okay. And so that's a  
13 dedicated system, I assume, just --

14 MR. PATEL: Correct. Just for this  
15 purpose.

16 MR. GARRIE: Do you guys differentiate  
17 between those? I notice in the answer there's a  
18 distinction between a couple milliseconds, a  
19 millisecond and other things, and that makes a  
20 little more sense.

21 Okay. So then how is the Activity Log data  
22 used by -- is Activity Log data used by systems in  
23 production beyond the AYI, DYI piece?

24 MR. PATEL: I do not believe so.

25 MR. GARRIE: Oh, this is an important

1     preface to tee it off.

2                 Generally speaking, I'm not looking for  
3     absolute. 95 percent is sufficient. Just tell me  
4     it's not absolute, but more likely than not. Right?  
5     I realize that there's a very loose definition of  
6     the control of some -- that underlie some of the  
7     systems and who can access and what gets put in  
8     them.

9                 MS. RING: Special Master Garrie, if I may  
10     on that topic.

11                I appreciate you saying that and just,  
12     again, for the comfort of the people who are here  
13     testifying today and so that they can be as helpful  
14     to you as possible, Mr. Patel isn't responsible for  
15     generating the Activity Log. So when he says, you  
16     know, he's not aware of its use in production, I  
17     just want him to feel comfortable to make the point  
18     I think you're trying to make, which is he doesn't  
19     know, but he's trying to tell you, you know, his  
20     understanding.

21                MR. GARRIE: Which is a lot better than  
22     mine.

23                MS. RING: What?

24                MR. GARRIE: Which is a lot better than  
25     mine.

1 MS. RING: Oh, your understanding?

2 MR. GARRIE: Yeah.

3 MS. RING: Yeah. I know the feeling.

4 Okay. Thank you.

5 MR. GARRIE: I mean, he knew it's in

6 [REDACTED] So he's ahead of me.

7 All right. So then I read the engineering  
8 process for adding data to the DYI, AYI, and  
9 Activity Log. I notice that there's a very  
10 prescriptive process, which there's quite a bit of  
11 thought.

12 But how do engineers determine if they need  
13 to add to DYI, AYI, and Activity Log; right? How  
14 does an engineer actually say, "Stick it here, not  
15 here. Put it there"?

16 MR. CLARK: I'll start. If anybody has  
17 more to add, they can.

18 That's, first of all, part of having our  
19 broad privacy program that we have. Every employee  
20 and every engineer is trained in privacy and have to  
21 recertify on a regular basis. Every product they  
22 build must go through a privacy review or a review  
23 for a variety of things.

24 As part of both that privacy program and  
25 the privacy review process is checks and evaluations

1 on data that's being added and where and how it  
2 should be mapped. And what we provided to you is as  
3 products are built or as products are being  
4 developed, the technical process for an engineer to  
5 actually add it to those product surfaces or to  
6 those product areas, AYI, Activity Log, and DYI.

7 MR. GARRIE: So then is it the engineer who  
8 makes the determinations of what goes in there?

9 MR. CLARK: The engineer might make a  
10 recommendation. It's ultimately the professionals  
11 in privacy review making the determination on the  
12 criteria for what -- whether that data already  
13 exists or whether that data should be added is part  
14 of that.

15 MR. GARRIE: That's why I was wondering.  
16 Right.

17 How does an engineer know if it's a  
18 duplicative cookie or not a duplicative cookie  
19 unless they are intimately familiar with the entire  
20 data structure of DYI, AYI? So I thought there may  
21 be a process for the engineer to determine or make  
22 that determination.

23 MR. CLARK: And that process is the privacy  
24 review process as part of that. They're provided  
25 all the context and information in the pages that

1     you have for the types of --

2             MR. GARRIE:   Yeah.   Yeah.

3             So that gives them the framework.   I'm just  
4     wondering -- they don't know the file; right?   So at  
5     some point the engineering team or the DYI group has  
6     got to be consulted, I assume, because you don't  
7     want to load the file with duplicative data.

8             MR. CLARK:   Correct.

9             MR. GARRIE:   Are there any other reasons  
10    why an engineer wouldn't include data in the DYI  
11    file then or the Activity Log?

12            MR. CLARK:   You know, there are a couple of  
13    examples.

14            I think one example might be that the data,  
15    for instance, being used may not be -- may not  
16    actually include the user or like actually know that  
17    it belongs to the user.   I think those are two  
18    different areas.   The data may not be associated  
19    with that user.   Like, for instance, somebody takes  
20    a picture of you and I and you upload it and never  
21    tag me.   There's a photo of me, but I don't know  
22    that that photo of me ever exists and so no one does  
23    since it was never tagged.   And, therefore, that  
24    wouldn't be.

25            The data may not be human readable.   There

1 might not actually be something to show.

2 MR. GARRIE: When you say human readable,  
3 all words are readable to me. So I always find  
4 human readable to be this ephemeral concept.

5 So for an engineer, human readable is  
6 something lawyers will never be able to read. Like  
7 you give a lawyer a for loop, a while loop, a  
8 do-while, variables, class files, whatever,  
9 that's -- I've had multiple lawyers tell me that is  
10 unreadable.

11 So how do you define human? When you're  
12 saying human readable, what are we saying, like?

13 MR. CLARK: Maybe a good concrete example  
14 might be like as things travel between multiple  
15 models or multiple classifiers, the model itself may  
16 attach a short-term or a transient -- a transient  
17 tag from an associative perspective that, you know,  
18 is still tied, but ultimately isn't even meaningful  
19 to the engineer that built the model.

20 MR. GARRIE: Is that the same for all data,  
21 mobile, everything, you know, TVs? I mean, you guys  
22 have all sorts of hardware; right? You have picture  
23 frames and TVs and all of this.

24 MR. CLARK: I don't think TVs, for the  
25 record.

1 MR. GARRIE: Well, it's accessible on a TV,  
2 I believe.

3 MR. CLARK: Sure.

4 MR. GARRIE: I guess you might treat that  
5 just as a normal screen with a web API. I don't  
6 know how you treat it. But like you have the  
7 frames, right, or you have a mobile device. Right.  
8 So there are 513 different versions of Android that  
9 are currently supported just in North America. So,  
10 I mean, everybody may have a different flavor.

11 So how are we defining machine readable,  
12 like? So are we saying, like, is a hash value  
13 non-mach- -- is that hash value, as an example, like  
14 you hash something, is that non-readable?

15 MR. PATEL: It depends what you're hashing.  
16 Sorry.

17 MR. CLARK: No. Please, Mr. Patel.

18 MR. GARRIE: So I'm just trying to figure  
19 it out because is it more the way the data is used  
20 or is it because a hash value is machine readable or  
21 does it include -- or hashes -- hashes would be  
22 non-DYI or DYI or how does that work?

23 MR. PATEL: So a hash of a picture, for  
24 example, photo DNA or something like that, we  
25 provide the original content that you uploaded to



1 Facebook. And so, therefore, any hash is a  
2 derivative of it in that sense.

3 MR. GARRIE: But there's metadata embedded  
4 in the photo. Like photos have -- you know, can  
5 have GPS. They can have -- depending on how fancy  
6 your camera is, it can have all sorts of things in  
7 it. But do you include that then in the DYI file?

8 MR. PATEL: I do not believe --

9 MR. GARRIE: The hash of a photo.

10 MR. PATEL: The hash of the photo we do not  
11 include in DYI. The original photo we do. The  
12 original photo also could be resized. Like if you  
13 upload a 16 meg photo, you know, we may not retain  
14 the 16 meg photo. We may retain a smaller version  
15 of it.

16 MR. GARRIE: I assume you'd resize it and  
17 put it in your blob and drop it into your photo  
18 system and archive it.

19 MR. PATEL: Correct.

20 MR. GARRIE: Yeah. So that's what I'm  
21 trying to figure out from the engineering side of  
22 the house what --

23 I guess what I'm really trying to figure  
24 out is how things get in the DYI and not in the DYI,  
25 but I think so then --

1 MR. CLARK: There are a few more things  
2 that don't go in DYI. I think there's -- there's  
3 some classifications both from like an integrity and  
4 security perspective that we don't put in there.

5 For instance, if we believe you're a bot,  
6 the process to classify that you're automated and  
7 not a real and authentic user might be one example  
8 of that.

9 MR. PATEL: And just going back to the  
10 hashing question.

11 The photo that you upload, when we provide  
12 that back to you, the hash of it, and especially  
13 photo DNA, should produce the same hash value. So  
14 it is effectively -- like, you can generate your own  
15 photo DNA hash from it.

16 MR. GARRIE: As long as you don't use -- as  
17 long as you embed data. It just depends. Yes. I  
18 agree. I do forensics work on photos.

19 MR. PATEL: Oh, okay.

20 MR. GARRIE: Perfect.

21 MR. PATEL: Yeah. I believe we strip the  
22 excess information if we do hash.

23 MR. GARRIE: So, yes. I believe you are  
24 correct having done forensics on Facebook photos  
25 myself.

1           Okay. So then I guess my question is, so  
2   when I read over the document, is there an  
3   independent from the engineer review of whether the  
4   data point is or is not added to the DYI, AYI or  
5   Activity Log?

6           MR. CLARK: That is the team of privacy  
7   review that does that review.

8           MR. GARRIE: And that's independent of the  
9   engineer?

10          MR. CLARK: That is correct.

11          MR. GARRIE: All right. That makes sense.  
12   All right.

13          And then a privacy XFN decision task, I  
14   notice that's the instrument that's needed to add  
15   data to the DYI; right? So how does the privacy  
16   XFN process work itself? Because that wasn't in  
17   there.

18          MR. CLARK: For the sake of clarity, that  
19   is the same team as the privacy review. The XFN  
20   stands for cross-functional and that is the prior  
21   name of the privacy review team.

22          MR. GARRIE: Privacy XFN process is the  
23   same as the one in the document?

24          MR. CLARK: Correct. The review.

25          MR. GARRIE: The review. But how does the

1 process then --

2 So when a decision task is needed to add  
3 the data, how does the privacy team process work?  
4 How does the privacy team do that work? Because  
5 that wasn't in there.

6 MR. CLARK: Sure. The initial product team  
7 building the product, part of what they're trained  
8 on on a regular and consistent basis and certify to,  
9 they create a path within privacy review in order to  
10 give the details of their product linked to code and  
11 get additional elements in order to have that review  
12 completed. As that review is completed, there are  
13 tasks that are identified out of that. Either  
14 remediations that were not done by the original team  
15 or the original team has done all of the things that  
16 they needed to do and then approval for those things  
17 to move forward.

18 MR. GARRIE: So as an engineer, does anyone  
19 actually make a -- sorry.

20 Does anyone actually make the developer,  
21 the engineer consult privacy and add the data to  
22 DYI? I'm not saying engineers aren't the most  
23 amazing people in the world. I'm just saying is  
24 there any checks and balances behind the process?

25 MR. CLARK: There are checks and balances

1 behind the process, which is part of why the task is  
2 created to track its completion and tie it to the  
3 disk that's actually deployed. It's part of why the  
4 privacy review exists. It's part of why like  
5 everybody is trained on it and agree to those  
6 obligations to even become an employee at Facebook  
7 or continue to be an employee of Facebook.

8 And so we take that program serious. So  
9 there is follow-up and there are checks and balances  
10 built into that.

11 THE COURT REPORTER: Mr. Garrie, I couldn't  
12 hear you.

13 MR. GARRIE: That's kind of funny. I was  
14 going to strike my comment anyway.

15 Off the record and we'll go back on.

16 (Discussion held off the record.)

17 MR. GARRIE: We'll go back on the record.

18 So then one other thing I didn't understand  
19 is what are the possible decisions? Because I tried  
20 to figure out what are the possible decisions that  
21 come out of the decision task; like approved,  
22 denied, incorrect. I was trying to figure that  
23 out.

24 MR. CLARK: So approved or denied or an  
25 incorrect.

1           Approved is all of the elements are  
2     outlined correctly and all of the both task and work  
3     is meeting the obligations identified in the privacy  
4     program.

5           A reject or the third option both means  
6     there is either some kind of remediation or this  
7     product cannot be released.

8           MR. GARRIE:   Okay.   What about schema and  
9     location or stuff like that?   I saw that referenced  
10    as well.

11          MR. CLARK:   As part of -- and I think  
12    having worked on systems for way too long, when I  
13    think of schema in traditional systems terms, it's  
14    here's the exact pointer to where it's sitting on  
15    blocks on disks.

16          Where schema is referred to here, rather,  
17    what it's pointing to is what -- we have a process  
18    we call [REDACTED]

19    [REDACTED] -- sorry -- [REDACTED]

20    [REDACTED]           And when we talk about referring

21    [REDACTED]

22    [REDACTED]

23    [REDACTED]

24    [REDACTED]

25          MR. GARRIE:   When you say give, I'm an

1 engineer, but you mean check and commit; right?

2 MR. CLARK: That is correct.

3 MR. GARRIE: Okay. Thank you.

4 MR. CLARK: I made the technical  
5 assumption.

6 MR. GARRIE: No. No. I agree. I just  
7 want everybody.

8 So check and commit is good and then polls  
9 requests. Okay. We're good.

10 So then [REDACTED]  
11 right?

12 MR. CLARK: That is my understanding, yes.

13 MR. GARRIE: All right. And then there's

14 [REDACTED] Okay. That makes  
15 sense to me. All right.

16 I welcome any other --

17 MR. CLARK: I don't know that that last  
18 statement is accurate.

19 MR. PATEL: I don't think the last  
20 statement is correct.

21 MR. CLARK: I wanted to correct that.  
22 There aren't things that live in DYI that  
23 aren't in DYI. DYI is the more complete set.

24 MR. MITCHELL: [REDACTED].

25 MR. GARRIE: Perfect. I get that. That

1 answered that question.

2 So then there was a lot of discussion about  
3 DYI. So I read in your information schema. It  
4 looks like there was a lot of resources devoted to  
5 that.

6 I guess my question to you gets more to --  
7 So question two is about what is in the  
8 DYI.

9 Is there anything else that anyone else  
10 wants to add to what we discussed so far about  
11 question one?

12 I got all my questions answered.  
13 Plaintiffs, now is the time to raise your hand if  
14 you have a question.

15 Going to the next question. All right.  
16 The second question. I read the answer and I have  
17 follow-on questions to that.

18 As I understand from the answer, the data  
19 is stored as raw, logged with event data in  
20 Facebook's data warehouse. Facebook then builds  
21 data pipelines from the data warehouse to  
22 production.

23 If I understood that right, what is the  
24 data warehouse called? I mean, you have dozens of  
25 data warehouses and, I assume, that's only a



1 fraction.

2 So in the response it says data warehouse,  
3 but we don't say which one.

4 MR. ZARASHAW: It is synonymous with Hive.

5 MR. GARRIE: And then when you say data  
6 pipelines, is that like a data pipe- -- like how are  
7 you defining a data pipeline?

8 MR. CLARK: Let's have Mayur, Mr. Patel,  
9 answer that.

10 MR. PATEL: Yeah. So a data pipeline is  
11 effectively processing -- a processing unit on top  
12 of the storage system.

13 So in our case Hive currently is synonymous  
14 for our storage as well as our processing pipeline.  
15 We also use other ones like Spark and Presto at the  
16 company.

17 MR. GARRIE: How are data pipelines  
18 themselves built?

19 MR. PATEL: Data pipelines are built by  
20 development.

21 MR. GARRIE: And particularly for the data  
22 warehouse and particularly for the raw logs and  
23 event data in Facebook's Hive. The data in the  
24 Hive, you build data pipelines. How are those data  
25 pipelines built for the Hive?

1 MR. PATEL: Yeah. So the most common one,  
2 and I'm just speaking for the generalized version of  
3 it, are built by data engineers or software  
4 engineers on top of Hive using these three  
5 processing pipelines that I talked about: Hive,  
6 Presto or Spark.

7 The language used for them is mostly SQL.  
8 So it's really easy to understand. I'm sorry. It's  
9 not standard SQL. It's not ANSI SQL.

10 MR. GARRIE: SQL. I can't remember. I've  
11 seen it.

12 But basically the Facebook SQL query  
13 language is used to build these pipelines? Is that  
14 what we're --

15 MR. PATEL: Yeah. Correct. It potentially  
16 is Presto SQL. It's not Facebook SQL, but it's not  
17 ANSI SQL. It's not the standard SQL that you know.

18 MR. ZARASHAW: Let me make a correction.

19 It is not Facebook SQL. There is an SQL  
20 which was -- that was a query used -- format used  
21 for the older versions of the developer platform and  
22 this is not it.

23 MR. GARRIE: So you're using Presto? You  
24 said Presto SQL?

25 MR. PATEL: Yeah. Primarily it's Presto

1 SQL. I believe Hive is fully deprecated and Spark  
2 is a future version of what we may use.

3 MR. GARRIE: How are the data pipelines  
4 executed?

5 MR. PATEL: They're executed by a scheduler  
6 called Dataswarm. So people write Python scripts  
7 that describe effectively what a unit of task is  
8 and what the interdependencies are. So an  
9 interdependency, for example, wait on table X to  
10 get data before you can process this meaningful  
11 task because if you don't have data in table X, you  
12 can't really go and query table X in your current  
13 task.

14 MR. GARRIE: Are they sequential? Because  
15 I looked at the ones of the list you gave me and  
16 some are repeated; right? So we'll talk about the  
17 exhibit itself.

18 But, I mean, the bottom line is they're run  
19 by Dataswarm effectively. It's a scheduler?

20 MR. PATEL: Correct.

21 MR. GARRIE: Okay. And then where is the  
22 data pipeline -- I might regret. Well, I don't know  
23 how else to ask this.

24 So where is the data pipeline's output  
25 stored?

1 MR. PATEL: The developers can choose where  
2 to store it. You know, most likely, they're storing  
3 it in another dataset to then be consumed by other  
4 data pipelines or then to be moved over somewhere  
5 else. But it's not uncommon for emailing stuff or,  
6 you know, like if you just want to email the output  
7 of a result or something like that.

8 MR. GARRIE: [REDACTED]  
[REDACTED]  
[REDACTED].

11 Are [REDACTED]  
[REDACTED] or something  
13 like that?

14 MR. PATEL: [REDACTED] [REDACTED]  
[REDACTED]

16 MR. GARRIE: And then so where is -- can  
17 they store it anywhere? Do they store it in the  
18 Hive? Is it stored in a subset of systems? Is  
19 it --

20 MR. PATEL: Primarily I would say most of  
21 it goes back into Hive.

22 MR. GARRIE: And then other parts of it, I  
23 assume, go to TAO, to other?

24 MR. PATEL: [REDACTED] [REDACTED]  
[REDACTED]

1 MR. GARRIE: Got it.

2 MR. PATEL: So we --

3 MR. GARRIE: Go ahead.

4 MR. PATEL: Go ahead.

5 MR. GARRIE: No. Go ahead.

6 MR. PATEL: [REDACTED]

7 [REDACTED].

8 MR. ZARASHAW: Pardon me for jumping in. I  
9 do have to clarify.

10 [REDACTED]

11 [REDACTED]

12 [REDACTED]

13 [REDACTED]

14 [REDACTED]

15 For example, for computing -- we do a very

16 [REDACTED]

17 [REDACTED]

18 [REDACTED]

19 [REDACTED]

20 [REDACTED]

21 There are certain cases where -- cases

22 [REDACTED]

23 [REDACTED]

24 MR. GARRIE: So then basically a dataset's  
25 a table in the Hive. For lack of a better, is that

1 the equivalent of what we're -- and then it's  
2 cheaper to do it there than to do it elsewhere, I  
3 guess, from a computation perspective.

4 MR. ZARASHAW: Exactly. So once in Hive  
5 per day or per week, and then redo it as many times  
6 as needed.

7 MR. GARRIE: So then returning back.  
8 When you said the dataset, Mr. Patel,  
9 that's a table. Is that what we're talking about?

10 MR. PATEL: Yeah. That's a table. That's  
11 effectively another Hive table.

12 MR. GARRIE: I technically get it.

13 MS. RING: Special Master Garrie, I'm  
14 sorry. I'd like to clarify something.

15 Mr. Patel has worked on these pipelines, I  
16 mean, for OFA, so in response to actually  
17 Question Number 2. So I feel like we're kind of  
18 talking in very general terms here, but this is just  
19 for OFA and the pipeline that was created from Hive  
20 into DYI for OFA. I just want to be clear about  
21 that and that they're talking about, you know,  
22 internal Facebook engineer developers and not  
23 third-party developers. So...

24 MR. GARRIE: Are third-party developers  
25 allowed to access the Hive?

1 MS. RING: What?

2 MR. GARRIE: I didn't think any third-party  
3 developers are allowed to touch the Hive.

4 MS. RING: They are not. That's the point  
5 I'm trying to clarify.

6 MR. GARRIE: Yeah. I'm assuming that  
7 anybody that's touching these dataset tables or data  
8 in the Hive is an authorized Facebook engineer doing  
9 some product-based specific activity for on-platform  
10 work or some analysis or something. Nothing to do  
11 with third parties at all. However, they're not  
12 allowed to touch it. They may have a contract and  
13 ask an engineer to do something and they may run  
14 something, but the Facebook engineer is running  
15 that. But I'm just trying to understand big picture  
16 for the languages we saw.

17 So then is there a list of the existing  
18 data pipelines? I mean, I got the five million, but  
19 is there an actual list?

20 MR. PATEL: Sorry. A list for all the  
21 pipelines at Facebook or the OFA pipeline?

22 MR. GARRIE: The OFA pipeline would be a  
23 good place to start.

24 MR. PATEL: Yeah. Yeah. We could  
25 potentially produce that.

1 MS. RING: I'm sorry. I have to clarify  
2 that.

3 Do you mean the name of the pipeline? I  
4 just want to be clear.

5 MR. GARRIE: No. I'll get to the points I  
6 want.

7 MS. RING: Okay.

8 MR. GARRIE: I haven't decided what I want  
9 yet or what I don't want. I just want to know what  
10 is possible.

11 MS. RING: Got it.

12 MR. GARRIE: Don't take any question here  
13 as a determination or direction of what I'm doing.  
14 I'm just learning here.

15 MS. RING: Yes.

16 MR. GARRIE: So then what about can you  
17 provide what data they can soon produce then for the  
18 data pipelines?

19 MR. PATEL: Oh, we could -- yeah. We could  
20 produce the SQL code.

21 MS. RING: Again, Mr. Patel is speaking  
22 only to the OFA platform. Okay?

23 MR. GARRIE: Yeah. I got it loud and  
24 clear.

25 MS. RING: I just want to be clear. I



1 mean, they're sworn. We need to be clear about what  
2 they are actually here to talk about.

3 So got it. I will shut up now. Thank you.

4 MR. GARRIE: I am assuming that we are  
5 speaking in very specific terms, but I assume  
6 someone else, since there is an engineer that did  
7 build the data pipeline, for whatever it is, can  
8 produce, that somewhere someone has the ability, it  
9 sounds like. There's tables, good news, and there's  
10 SQL, well, Presto SQL, and you have input and you  
11 have output and there's an engineering resource that  
12 did create that whatever -- whatever they're running  
13 in sequence that run are managed by the Dataswarm.  
14 Someone had to schedule it, approve it, and put it  
15 in and I get that.

16 MR. CLARK: And there's such a long list  
17 because you can create a pipeline for temporary  
18 data access. It's not generating new data. It's  
19 only using existing data. So you can do that for,  
20 as you mentioned earlier, research, trying to help  
21 diagnose something that may be going on or like the  
22 example that Mr. Patel is referring to with OFA is a  
23 great example operationally to look at.

24 MR. GARRIE: They can literally do it for  
25 whatever.

1 MR. CLARK: Correct.

2 MR. GARRIE: It's like basically a data  
3 junkyard.

4 MR. CLARK: Yes, it is, but it's also very  
5 governed under the privacy program where the  
6 biggest part of the training is what access to  
7 data you have as an employee and how you are  
8 supposed to treat and handle that access to data as  
9 well. So --

10 MR. GARRIE: I agree. There's a privacy  
11 construct. Like you said, all the engineers get  
12 trained and when they come on they sign these  
13 compliance, they got to recertify. They got lots of  
14 paperwork regarding the process on data privacy  
15 management. I'm just trying to understand. So  
16 there's not -- I'm not assuming anything bad. I'm  
17 just trying to understand where it all sits and how  
18 complicated these tasks are.

19 So then jumping in. Category 1,  
20 on-platform. So this is user-provided data. Is all  
21 user-provided data included in the DYI tool? I  
22 mean, user provided. So I'll be very clear.

23 MR. CLARK: To the best of my knowledge,  
24 yes, the user-provided data is.

25 MR. GARRIE: Okay. Are there individual

1 fields from schemas that are not included in the DYI  
2 file?

3 MR. CLARK: Would you repeat that question?

4 MR. GARRIE: Are there individual fields  
5 from schemas that are not included in the DYI file?

6 MR. CLARK: I do not know.

7 MR. MITCHELL: I'm just clarifying.

8 Are you asking if we have a table that has  
9 a series of columns about a particular user or are  
10 we excluding a particular column? Is that what  
11 you're trying to get at?

12 MR. GARRIE: A column. A value. I mean,  
13 you guys call them -- yeah. In essence, like, I  
14 don't know how -- yeah, a schema. Like a schema can  
15 have columns, rows, depending on how you define the  
16 object; right? I mean, it's all relative to the  
17 design of the object and how you put the tables and  
18 link the tables and everything else.

19 But the question is, they're fields that  
20 don't make it in. I mean, so here's my -- my --  
21 well, let's actually differentiate because my next  
22 question is about cold tables.

23 I'm talking about schemas. So field values  
24 versus tables; right?

25 So because are there -- so like a schema --

1 a field in a schema is like a -- I don't know how  
2 you guys refer to tables and datas and rows and  
3 values. But an individual entry, okay, from that,  
4 are those things not included in the DYI file?

5 MR. CLARK: I believe if we look at the  
6 process for an engineer to add --

7 I apologize. I can't just scroll through  
8 it.

9 MR. GARRIE: I know.

10 MR. CLARK: So part of the mapping in the  
11 engineer's responsibility and what goes through the  
12 privacy review is mapping the complete -- mapping  
13 the data that the user provided to that output for  
14 DYI output.

15 I think what you're asking, and this is why  
16 I'm asking from a clarity perspective, is --

17 MR. GARRIE: It's okay.

18 MR. CLARK: -- if there's something that's  
19 stored raw in that row, that there's some other set  
20 of data that may not go with that.

21 MR. GARRIE: Yeah. At a field level and  
22 then at a table level.

23 MR. CLARK: I'm not aware of any, but my  
24 better answer would be I don't know.

25 MR. GARRIE: Okay. We'll stick with I

1 don't know and we can double-check because that  
2 will be -- because there's an engineering review  
3 process and I assume that -- well, I don't know --  
4 to Ben's point, to Mr. Mitchell's point, I don't  
5 know how you refer to it, but I'm talking about a  
6 schema being a user's an object. My eye color's  
7 hazel, my height, all of those things, right,  
8 however you put, like, the eye color or the  
9 circumference of my eye would be a field value and  
10 then maybe the whole table describing my eye, just  
11 as an example; right? So --

12 MR. MITCHELL: Yes. So as I called out in  
13 my declaration, like there are certain types of  
14 things that we don't include deliberately. So it is  
15 possible that there is a column that contains a  
16 trade secret and that column -- and so we would not  
17 include that in DYI. Right.

18 MR. GARRIE: Right.

19 MR. MITCHELL: And we're not going to  
20 discuss on this call what those trade secrets might  
21 be obviously, but like I think -- so it is certainly  
22 possible that there are things that are excluded at  
23 the field-by-field level.

24 MR. GARRIE: All right. And so that's not  
25 user provided.

1 MR. MITCHELL: Fair enough.

2 MR. GARRIE: Users aren't providing trade  
3 secrets. I'm only talking about user provided.

4 MR. PATEL: Yeah. We provide all user  
5 provided. If you wrote in -- yeah. If you -- if we  
6 had an option for eye color, we would provide that  
7 in DYI and the user chose that.

8 MR. GARRIE: Yeah. We'll get to the other  
9 ones in a second.

10 MR. MITCHELL: Okay. Yeah.

11 MR. GARRIE: Because you guys  
12 differentiated it in your answers, user provided,  
13 non-user provided. So I'm just going through it  
14 that way rather than -- yeah.

15 MR. MITCHELL: Okay.

16 MR. GARRIE: And everything I give comes  
17 in. Okay.

18 Counsel Weaver, what question do you have?

19 MS. WEAVER: I would be interested in the  
20 definition of user provided and this is the reason.

21 For example, if I went on a website and off  
22 platform, but on my Facebook log in and put  
23 something in a shopping cart, but then never  
24 purchased it, that is information that we do not see  
25 in the Download It Yourself tool that our experts

1 have said is certainly some information that  
2 Facebook collects, but maybe that's not determined  
3 to be user provided.

4 MR. GARRIE: You don't have to answer the  
5 questions directly.

6 How are you -- I thought you provided the  
7 definition of user-provided data, but how are you  
8 defining user-provided data? Is it any different  
9 than what was provided?

10 MR. CLARK: I'm going to refer to the  
11 definition that we included where it was directly  
12 user provided from the first-party platform itself,  
13 from Facebook.

14 MR. GARRIE: That's what I understood it  
15 as well. I was using the definition that you  
16 provided.

17 So that would not be defined,  
18 Counsel Weaver, as user-provided data.

19 MS. WEAVER: Right. I mean, and let me be  
20 clear.

21 With first-party and third-party cookies,  
22 arguably Facebook could have said if they use a  
23 Facebook login, that's a first-party data point, but  
24 I hear their definition here is saying it is not.  
25 Is that correct?

1 MR. GARRIE: Well, you can read -- I mean,  
2 if you read their answer, they have the written  
3 definition in the engineering.

4 MS. WEAVER: Okay.

5 MR. GARRIE: And I don't know what page  
6 it's on, but I'm certain I've read it.

7 MS. WEAVER: I've read it. It's just  
8 unclear to me, but I understand the position they're  
9 taking. Okay.

10 MS. RING: I'm sorry to interject, but the  
11 question that Counsel Weaver asked, I mean, what  
12 you just described there and, you know, Mr. Patel  
13 is here. He worked on the OFA. It's just -- it's  
14 not in the first category, but it's in the off  
15 Facebook.

16 MR. GARRIE: I know. We're going to get to  
17 it. That's why I'm not worried about it.

18 MS. RING: Okay. I just didn't want to  
19 leave the impression it's not in the DYI file. It's  
20 just not called user provided.

21 MR. GARRIE: Yeah. I get it. I'm fully on  
22 board. You guys worked hard at this. I'm just  
23 going through methodically the different pieces.  
24 But her question was specifically about how user  
25 provided is defined and we're only talking about



1 on-platform; right? So we're still in Category 1,  
2 on-platform. We haven't even departed.

3 So observed data. Okay. So that's the  
4 next bucket here. How is observed data collected  
5 from the on-platform activity?

6 MR. CLARK: Yeah. I think we tried to  
7 outline that like exactly with concrete examples and  
8 exactly for what it is.

9 It includes interactions like clicks, what  
10 profile, pages, groups, and events the users have  
11 visited, the usage data, the device data, for  
12 instance, data about the user's activity level,  
13 advertisers with which the users interacted, IP  
14 addresses, IP address, where they have users that a  
15 user has chosen to see less or see first in news  
16 feed, time spent watching from a page of the watched  
17 product, people whose profile a user has visited,  
18 last location, last active time.

19 I could keep going through lists, but we  
20 tried to provide like a good set of examples of what  
21 are observed.

22 MR. GARRIE: All right. So I'll get to  
23 Counsel Weaver's question in a few minutes.

24 But what I don't understand is how are you  
25 collecting it? Are you using like API? Like I

1 guess what I'm trying to understand is, there's a  
2 lot of data you can observe. Right. There has to  
3 be -- is there a table schema? I mean, is there  
4 some -- I mean, are you collecting it into tables?  
5 I'm just trying to understand how is it actually --

6 I appreciate the examples. They were  
7 helpful.

8 What I'm trying to understand is, how is it  
9 actually being collected from an on-platform  
10 activity? Is it being stored to, like, the -- like,  
11 how is it actually working?

12 MR. CLARK: Through the series of examples,  
13 they're all not stored in one place. Like, they are  
14 part of the product.

15 So [REDACTED]

[REDACTED]

[REDACTED].

18 Examples like [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED].

22 MR. GARRIE: But there's no unified  
23 perspective of where all of the user observed data?  
24 Like, there's no observed data table set or object  
25 or schema or -- because I guess by product,

1 different data is -- by different use of the  
2 product, different data is observed. But a user,  
3 based on the products they use, there's -- is it in  
4 a schema?

5 I mean, how is the observed data being  
6 collected? Is it going into a schema, in a table,  
7 and then associated or on the platform or is it by  
8 product? Is that how that's working?

9 MR. CLARK: It's a combination.

10 MR. PATEL: Yeah. I was just going to --

11 MR. CLARK: Mr. Patel.

12 MR. PATEL: Yeah. Like, it depends on the  
13 product. I think you nailed it right there.

14 There may be cases where it's stored in the  
15 data warehouse. There may be cases where it's  
16 stored in our graph databases.

17 And what we do is the ones that are -- like  
18 the ones that we have access to, like that's the  
19 whole [REDACTED] question that you had, we're able  
20 to produce those in Activity Log. While we're also  
21 producing those in Activity Log, we can produce  
22 those in DYI.

23 So some examples of it, you know, [REDACTED]  
[REDACTED]  
[REDACTED]

1 [REDACTED] for eventually to be downloaded in DYI.

2 MR. GARRIE: Okay. So then is the observed  
3 data generated from the warehouse via data  
4 pipelines? Is that how that's working on a  
5 product-by-product basis or --

6 MR. PATEL: I don't know the intrinsic,  
7 like, pipeline of the pages you may use. I don't  
8 know if Eugene knows it. But I would suspect that

9 [REDACTED]

[REDACTED] [REDACTED]

[REDACTED] [REDACTED]

[REDACTED] [REDACTED]

13 MR. ZARASHAW: I could try to help you out.

14 There are a number of different  
15 infrastructure pieces that can be used for this  
16 purpose. Each product team decides at the time  
17 they're implementing the product what outcome  
18 they're trying to achieve and, therefore, which  
19 pieces of the infrastructure can best match the  
20 outcome.

21 So if I could use a couple examples. One  
22 would be, let's say, you like a page on Facebook.  
23 That specific example happens to immediately go  
24 into TAO, but the fact that you liked the page is  
25 recorded in TAO as your relationship to the page

1 because in order for this to work, you need to also  
2 be able to see that you liked that page next time  
3 you reload that page, which means there has to be  
4 immediate real-time access to the fact that you  
5 liked it available through the web content.

6 There are other cases where we might have  
7 observed data like, let's say, the IP address a  
8 person is visiting from. We might not even use it  
9 and it ends up in Hive as a log and that might be  
10 the end of it. It just sits in Hive.

11 So it really depends on somebody coming up  
12 with the scenario for how the data is to be used  
13 before anything is done with it.

14 MR. GARRIE: Based on the product usage  
15 effectively, and then from that you can find the  
16 observed data by the product that's used by the  
17 user.

18 MR. ZARASHAW: Exactly. And how that  
19 product team chose to implement that specific  
20 observed data.

21 MR. GARRIE: And then observed data sounds  
22 like can come from pretty much anywhere. I couldn't  
23 tell from the examples.

24 From what I could tell, I'm just trying to  
25 list out where observed data is collected from about

1 a user, and it looks like it comes from lit- --  
2 like, based on my examples, you cover pretty much  
3 everything.

4 MR. CLARK: Yes. That goes back to like  
5 that's why we walk through those in privacy review.  
6 And even just looking through my own, like, I have a  
7 very, very consistent observation. Like, that is  
8 the intention of the observed data.

9 MR. GARRIE: But observed data can be  
10 stored based on how the product is built and  
11 architected. It may be stored in the Hive. It may  
12 be stored in TAO. It may be stored in [REDACTED]  
13 Like, I don't think it's directly to [REDACTED] but  
14 somewhere it ends up in [REDACTED] somehow that goes  
15 to the social graph or MySQL or whatever the  
16 engineer and the privacy team thinks is the  
17 appropriate resource and it controls for that  
18 observed data to treat it accordingly and store it  
19 accordingly.

20 So then how much -- so then the other thing  
21 I'm trying to figure out is how much observed data  
22 is stored? For example, like how many page visited  
23 are stored, how many group visits are stored, how  
24 many devices is the data stored about? You know,  
25 sort of, is there any limit?

1 MR. CLARK: Could you clarify what you mean  
2 by that question?

3 MR. GARRIE: So observed data, you can  
4 observe all sorts of things; right? The question  
5 is, how much observed data is actually stored about  
6 a user and how many pages -- do you know how many  
7 pages are stored? Like, is that stored for how  
8 like -- like, is there anybody that would know or is  
9 it on a per-product basis?

10 MR. CLARK: It's -- because of our scale,  
11 that's going to be continually dynamic. And as  
12 products get built, as products get deprecated,  
13 that is going to be another thing that constantly  
14 changes.

15 And are you looking for the quantity of  
16 like how many clicks were tracked or are you looking  
17 for the number of places?

18 MR. GARRIE: What I'm looking for is, so  
19 let's say you deprecate how many clicks are tracked.  
20 Right. How many -- like how many -- I guess how  
21 much do you store, like for a user is actually  
22 stored, but it sounds like that's determined by the  
23 product and then the product defines the amount of  
24 observed data and the timeline it will store and  
25 capture it for. Is that it?

1 MR. CLARK: That is accurate and correct.  
2 The combination of the product team and the privacy  
3 review team.

4 MR. GARRIE: So then my next question is,  
5 I'm reading over the DYI and AYI. So then -- well,  
6 I'll give you a practical example.

7 For a while I was pretty good at losing  
8 phones, like masterful actually. So, I mean,  
9 believe it or not, I don't know if I'm on my 20th,  
10 but let's just pretend I am. I'm on my 20th phone  
11 and do you have all of them or just the last or is  
12 that product specific?

13 MR. CLARK: I can share, based on my own  
14 AYI file, I observed multiple devices that I've  
15 switched over time.

16 MR. GARRIE: But I guess it depends on the  
17 product; right? The mobile product defines how long  
18 they're going to capture that observed data, and  
19 then that then goes through the engineering review  
20 process, then says this belongs in the DYI file. So  
21 it's sort of driven that way rather than some de  
22 facto rule about we keep this much for this long per  
23 this. It's per product and if you deprecate it or  
24 you don't deprecate it. I mean, people have Palm  
25 devices still, so...



1 MR. CLARK: There are defined retention  
2 periods that are part of that privacy program and  
3 part of that as well.

4 MR. GARRIE: So there are defined retention  
5 programs for observed data as well?

6 MR. CLARK: I cannot specifically say for  
7 observed data versus other kinds of data. I just  
8 know there's those programs.

9 MR. GARRIE: Per product. So the analysis  
10 is done by product and then the data on them by  
11 product have policies that govern those products.

12 MR. CLARK: I could not speak directly to  
13 that, but --

14 MR. GARRIE: That's okay.

15 MR. CLARK: -- that's the general concept,  
16 yeah.

17 MR. GARRIE: I understand.

18 And so then I guess my other question I  
19 was trying to figure out, is all observed data  
20 included about a user included in the DYI file, the  
21 DYI tool? Sorry. You called it a tool. The DYI  
22 tool.

23 MS. RING: I think Mike is locked up.

24 MR. GARRIE: My kids do that when we're in  
25 school sometimes. When they're having remote

1 school, they pretend to freeze like this.

2 MR. CLARK: I didn't know if that was you  
3 freezing or me freezing.

4 MR. GARRIE: That's fine. Okay.

5 So the question I asked is, is all observed  
6 data included in the DYI tool? I couldn't tell.

7 MR. CLARK: There are -- you know, I can  
8 talk to some principles on why we would not include  
9 something. And I think an example of something that  
10 might be observed might be some behavior that, you  
11 know, from an integrity perspective, that you  
12 started the process of flagging something as  
13 objectionable or violating our terms and then  
14 stopped flagging that along the way.

15 There is some observation that, you know,  
16 you started the process of flagging something or  
17 reporting something; but when you didn't complete  
18 it, that wouldn't be in there.

19 MR. GARRIE: There's data in the observed  
20 data bucket that's not in the DYI file. It will  
21 vary based on product basically is what you're  
22 saying.

23 MR. CLARK: Based on product and based on,  
24 like, the set of -- you know, it's we can associate  
25 directly to you or, you know, it's something for

1 security or an integrity perspective.

2 MR. GARRIE: So is there any way to know  
3 what is excluded?

4 MR. CLARK: The items that meet those  
5 criteria, but I don't have an answer beyond that.

6 MR. GARRIE: And then because I'm just  
7 wondering like if there are individual fields that  
8 are excluded from the schemas or there are entire  
9 tables that are not included for observed data. Not  
10 trade secrets. Observed user data.

11 MR. PATEL: We've included the most  
12 meaningful ones that we believe that users want.  
13 Like, there's obviously product experiences that  
14 we've kind of excluded out.

15 For example, if you -- we have a product  
16 like Marketplace. If you open Marketplace for the  
17 first time, we may show you an introduction to  
18 Marketplace and, like, what all the -- this is  
19 hypothetical. So like an intro screen to like what  
20 all the buttons on Marketplace are I believe we  
21 provide when you open Marketplace, but we wouldn't  
22 provide you with the information that you went  
23 through; the tour, for example. That seems to be  
24 meaningless. It's also duplicative in that sense.  
25 But that is for the product experience

1 specifically.

2 MR. GARRIE: So then there's no delineated  
3 way to determine what individual fields in tables  
4 from the observed data per product have been  
5 excluded from the DYI file as we sit here today  
6 unless you go to each product and figure it out;  
7 right?

8 MR. PATEL: Yeah. Correct.

9 MR. GARRIE: Again, no decisions are being  
10 made here just to head off any questions. I'm just  
11 asking questions.

12 MS. RING: Special Master Garrie, I was  
13 going to suggest maybe a break, but I was also  
14 wanting to maybe let you get through. I know you're  
15 going category by category. So I don't want to  
16 interrupt that flow.

17 MR. GARRIE: Counsel Weaver, you have a  
18 question.

19 We are going to take a break. Well, first  
20 we're going to take Counsel Weaver's question.

21 MS. RING: Oh, okay.

22 MR. GARRIE: Then we're going to take a  
23 break.

24 Just so we're all clear, the next category  
25 is off-platform, Category 2. So we want to keep our

1 questions limited to Category 1, on-platform.

2 That's what we're talking about.

3 Counsel Weaver.

4 MS. WEAVER: Understood.

5 For on-platform activity just  
6 descriptively, would drafts of posts or drafts of  
7 messages be included?

8 MR. CLARK: I would not know.

9 MR. GARRIE: Wait. I was -- I was going to  
10 say I -- no. It's all right, Mr. Clark. I get  
11 engineers want to answer questions. I'm the same  
12 way. When my wife asks me questions, I just  
13 automatically answer them. Well, for many  
14 additional reasons, which I just answer. I guess  
15 it's the fastest way to avoid pain.

16 But the point being is, I was interested in  
17 that question as well. So if you wouldn't mind  
18 answering, Mr. Clark, the question or anybody else.

19 MR. CLARK: With that specificity, I would  
20 not know without needing to do further research.

21 MR. GARRIE: I mean, it's a  
22 product-by-product basis; right? And there's a ton  
23 of deprecated products and I assume, based on the  
24 size of the datas that we're looking at, there's a  
25 lot of it and it's compressed and everything else.

1           So they would need to go back and look on a  
2     product-by-product basis.

3           MS. WEAVER: I would just state  
4     Facebook Messenger is Facebook's own product and I  
5     think that would be helpful.

6           MR. GARRIE: We're only talking about  
7     Facebook, though.

8           MS. WEAVER: Okay. Fine.

9           And then the second question is the time  
10    period overall, Special Master Garrie. It would  
11    be -- I don't know if we're talking present when  
12    you're receiving your answers. Maybe you are, maybe  
13    the witnesses aren't, but it does seem like there's  
14    some ambiguity.

15          MR. GARRIE: I'm assuming we're only  
16    talking as of today, not anything deprecated in the  
17    past.

18          Is that an inaccurate assumption, Mr. Clark  
19    or Mr. Mitchell or Mr. Patel or Mr. Zarashaw? We're  
20    talking about the way it exists today, not the way  
21    it existed in the past; right?

22          MR. CLARK: That is accurate and correct.

23          MR. ZARASHAW: Correct.

24          MR. PATEL: Correct.

25          MR. GARRIE: You got to say yes,

1 Mr. Mitchell, because they won't transcribe it  
2 otherwise.

3 MR. MITCHELL: Yes.

4 MR. GARRIE: I give thumbs up, too, and  
5 then people get confused when they read the  
6 transcript.

7 All right. So we're going to take a  
8 five-minute break, if that's okay with everybody,  
9 and then we'll adjourn. I appreciate it.

10 Is that all right? Do you guys need more  
11 than five minutes? Do you want ten minutes?

12 MS. RING: Maybe ten just because by the  
13 time we get, whatever, and people can go get a drink  
14 or whatever.

15 MR. GARRIE: Yeah. Hopefully the JAMS  
16 moderator is here and everybody can go to their own  
17 rooms. If not, that means I got to assign  
18 everybody. I'll just open all rooms, then I'll  
19 close them all.

20 THE COURT REPORTER: Are we off the record?

21 MR. GARRIE: Go off the record and strike  
22 the last three sentences.

23 MS. RING: I'm not sure you can do that,  
24 Mr. Garrie.

25 (Recess taken.)

1 (Off the record at 11:36 a.m. Back on the  
2 record at 11:54 a.m.)

3 MR. GARRIE: We'll go back on the record  
4 and get started. Back on the record.

5 We finished Category 1, on-platform  
6 observed data, and we're now going to Category 2,  
7 off-platform.

8 So my question here, and I did my best to  
9 try to understand the information that was provided  
10 in the prior information and all of that so I keep  
11 it efficient, what type of third parties provide  
12 off-platform activities? Like, I couldn't figure  
13 out based on my look through the DYI file.

14 So I want to know what type of third  
15 parties provide off-platform activities?

16 MR. CLARK: Third-party developers. So  
17 that could be third-party developers using the SDK.  
18 It could be third-party developers integrating with  
19 the ads API.

20 MR. PATEL: Mr. Garrie, it's specifically  
21 developers that have integrated our business tools.  
22 So it's not any app. It is basically developers  
23 have chosen to integrate the business tools. Like  
24 app SDK, Pixel is another example, Facebook social  
25 plug-ins, CAPI is another one, conversion API --



1 apologies for that acronym -- and the ads API.

2 MR. GARRIE: Is there a list of all of  
3 the -- you gave me a list of products, but what I  
4 couldn't figure out is, so it's integrating the  
5 integrates to Facebook's app business tool API.  
6 Does that include deprecated?

7 MR. PATEL: What do you mean by deprecated?

8 MR. GARRIE: I can walk you through the  
9 license --

10 MR. PATEL: What do --

11 MR. GARRIE: -- (inaudible) community.

12 MR. PATEL: Sorry?

13 MR. GARRIE: Ones that you no longer want  
14 to offer to the development community, but people  
15 have built against and have invested substantial  
16 resources and time to integrate into, but you don't  
17 want to make it available to other new people, but  
18 you keep it around for deprecated. Like Windows 95  
19 still works. It's completely unsupported and  
20 there's no patching for it, but you can still turn  
21 on a Windows 95 computer and still use Windows.

22 MR. PATEL: Maybe I'll walk you through how  
23 a developer integrates the Facebook business tool  
24 because I don't understand the analogy with  
25 Windows 95.

1 MR. GARRIE: All right. So basically I'll  
2 give you a specific example based on what I read  
3 from what you guys gave me.

4 Say I was Netflix five years ago and I  
5 integrate into a now no longer available Facebook  
6 business tool. Five years ago there were business  
7 tools available that are no longer available today,  
8 but I've had substantial resources to integrate  
9 that. You've turned off -- you no longer offer  
10 that service to the marketplace, but you have an  
11 agreement with Netflix that says we're going to  
12 provide this SDK to integrate. You agree to do  
13 this, we agree to do that and we will continue to  
14 support it, but you no longer offer that as a  
15 developer SDK for any net new people or companies.

16 I mean, I can cite interrogatories. You  
17 have literally lots of them.

18 MR. PATEL: In your question, would that  
19 database be in OFA?

20 MR. GARRIE: Well, my question is, do you  
21 still -- do they -- when you said business tools,  
22 my question was is it deprecated business? Is it  
23 just business tools today or is it business tools  
24 since 2007?

25 MR. PATEL: It would be --

1 MR. ZARASHAW: There are a number of --  
2 Go ahead, Mayur.

3 MR. PATEL: No. Go ahead, Eugene.

4 MR. ZARASHAW: There are a number of  
5 platform APIs we have fully deprecated over the  
6 years where we -- one of the changes in platform  
7 1.0 to 2.0 was the introduction of a deprecation  
8 process where, for a lot of the APIs, they truly  
9 stopped working after a certain number of days after  
10 being deprecated and, yes, it meant that all persons  
11 have to update their software to deal properly with  
12 it, but we don't leave every single integration in  
13 API running in perpetuity.

14 MR. GARRIE: My question is, when you say  
15 Facebook business apps that they integrate to the  
16 SDK, do you have a list of which apps? Like,  
17 because those were available in the marketplace, but  
18 there are ones that are no longer available that, I  
19 assume, you still have agreements with developer  
20 third parties where you have to continue to support  
21 it at some level.

22 MR. ZARASHAW: We do not have that list  
23 right now. With a significant amount of work it  
24 could be provided, but it would number in the  
25 millions. And a large number of these also would

1 not have specific contracts and we might not even  
2 have an understanding of exactly how they use the  
3 business products because for certain features of  
4 the developer platform there is just a very simple  
5 sign-up process that does not require review unless  
6 they want access to the features that do require  
7 review.

8 So an example would be if a website wanted  
9 to put a light button there, anyone is free to  
10 register an app ID on Facebook and put a light  
11 button on their website and that would be an  
12 integration or an app.

13 MR. GARRIE: That's how I understood it as  
14 well.

15 So my question then becomes, at a practical  
16 level, third-party -- and I guess we'll get to this.

17 You know, one of the business tools listed  
18 is off-line conversions. So that's what I have a  
19 particular question about because I didn't  
20 understand what that tool actually does.

21 But before we get to that, I guess there's  
22 no easy way -- well, you at least have --

23 So anybody that integrates with Facebook  
24 APIs or SDKs as a third party may be providing  
25 off-platform activity back in some level?

1 MR. ZARASHAW: Yes.

2 MR. GARRIE: Just at a high level, where is  
3 data collected via the business tool stored?

4 MR. ZARASHAW: So that goes back to the  
5 pattern from earlier, which is the business tools  
6 are an external facing set of interfaces, i.e., the  
7 Graph API and some other tools, that front a number  
8 of products.

9 So depending on the product represented,  
10 the storage could be different. It's however that  
11 product chose to store the data that's relevant.

12 MR. GARRIE: The data that's collected via  
13 business tools, is it extracted and processed all on  
14 the data pipeline from the warehouse? Is that how  
15 that works?

16 MR. ZARASHAW: Not necessarily. Again, it  
17 would depend on the product. So some specific  
18 products may use production databases for the right  
19 immediately. Some products may use data warehouses.  
20 Some products may have product specific data  
21 storage.

22 MR. GARRIE: And these products will define  
23 how the data is extracted and processed and where it  
24 gets stored?

25 MR. ZARASHAW: Yes.

1 MR. GARRIE: Okay. And then so I  
2 understand the off-line conversion, what is that,  
3 business tool?

4 MR. PATEL: The off-line conversion  
5 business tool -- and I'm not an admin engineer, but  
6 I will give you a high-level summary of what it  
7 is -- is when an advertiser advertises on our  
8 platform and wants to understand store sales, so not  
9 specifically website sales that they may have, they  
10 can upload information about purchases that happened  
11 in their store with specific PII and we would  
12 attribute it back to users that may have seen their  
13 ad in order to provide them with a return on ad  
14 spend.

15 MR. GARRIE: Yeah. So you're basically  
16 giving them the Holy Grail of walk in the store, see  
17 a digital ad, can I connect it to a sale.

18 Okay. And then is there documents -- I'm  
19 trying to think of a better --

20 Each business Facebook business tool, does  
21 it have its own documentation about -- about what  
22 data is collected? And not just today, but for  
23 whatever is -- not only what you offer today, but  
24 what has been offered; what is still available, not  
25 necessarily offered.

1           So I'm distinguishing between offered and  
2     available. Offered means we promote this in the  
3     marketplace. Available means you know it exists, we  
4     got permission to use it, it's available to you. So  
5     I'm distinguishing between the two.

6           So my question is, for any business  
7     product, is there documentation -- for each business  
8     product that is available, is there documentation  
9     about where the data is collected and stored?

10          MR. CLARK: So tying to that question, for  
11     each of the business products -- and I hope this is  
12     answering your question, Special Master Garrie.

13          For each of the business product  
14     integrations, we do have very extensive  
15     documentation for those business products that  
16     instruct the engineer or developer or third party on  
17     how to integrate in detail on the data that is part  
18     of that integration and how to use that and all of  
19     that is public documentation.

20          MR. GARRIE: Yeah. I looked at the  
21     engineering website.

22          So is that for anything that's available or  
23     only what is offered?

24          MR. CLARK: So we're talking specifically  
25     about the business products integrations here. We

1 have our third-party developer platform, which you  
2 probably saw that documentation as well.

3 MR. GARRIE: They're different, though.  
4 They're completely different.

5 MR. CLARK: They are completely different.  
6 That is correct.

7 MR. GARRIE: Yeah. I mean, called  
8 Search and State. They're different objectives  
9 and, I mean, both code. We can both agree they're  
10 software code, so -- and APIs, but they're  
11 different. Right.

12 So the business products that are  
13 available, is that true for all of them that are  
14 available or only --

15 MR. CLARK: So for all of those programs,  
16 that documentation is available. That is not to  
17 say that there are not additional APIs that are  
18 developed as part of this, but those -- like, for  
19 instance, I think a good example of this might be  
20 the measurements API, which is available to a  
21 limited number of folks to help with ad measurement,  
22 for instance. But those, any APIs that are  
23 developed, I'm going to broken record it a little  
24 bit, like are part of both the privacy review  
25 program in addition to a special API XFN that goes



1 through and reviews any additions or changes how  
2 people integrate any data, read or write, that is  
3 shared to make sure that we have proper consent and  
4 everything else in addition to the privacy review,  
5 which would make sure that that data would end up in  
6 DYI.

7 MR. GARRIE: I'm just thinking about the  
8 answer.

9 MR. CLARK: Oh, yeah.

10 MR. GARRIE: Are there private APIs for  
11 business?

12 MR. CLARK: I am aware of the business  
13 products API platform, but I -- I personally don't  
14 know of any. That doesn't mean that they don't  
15 exist.

16 MR. GARRIE: They would have to be  
17 documented at some point.

18 MR. CLARK: Correct.

19 MR. GARRIE: There would be some process.

20 Who has the -- so then who has the master  
21 list, so to speak? Who runs the API, like the XFN?  
22 So who runs the API XFN?

23 MR. CLARK: The API XFN is run by our  
24 developer platform team and they do the reviews of  
25 those.

1 MR. GARRIE: They would have a list, I  
2 assume, of the work they do; right?

3 MR. CLARK: They -- like, with significant  
4 effort, we could get that list from either them or  
5 from how we track our third-party obligations.  
6 Like, that list will exist in some way somewhere.

7 MR. GARRIE: I'm not saying I want any  
8 list, just to repeat myself again. I'm just trying  
9 to figure it out.

10 And then for the off-line conversion tools,  
11 I guess there's a specific person at Facebook for  
12 each of these tools that would tell me where --  
13 what's collected, where it's stored, where it's  
14 extracted, and where it's processed. Is that an  
15 accurate assumption or can somebody here answer the  
16 question?

17 MR. CLARK: For that full life cycle, that  
18 would be product by product. I believe for some of  
19 the OFA components, which is the acronym we use for  
20 off-line activity, that Mayur can speak to elements  
21 of the life cycle of that data and how we present  
22 that back.

23 MR. GARRIE: So then for off-line -- for  
24 off-line conversion tools, where is that data  
25 stored?

1 MR. PATEL: Yeah. The data that's pulled  
2 or basically effectively brought into the OFA  
3 transfer, which then would be brought into DYI, is  
4 stored in Hive. We store it on a [REDACTED]  
5 [REDACTED] We then  
6 [REDACTED] as well as  
7 [REDACTED] and then [REDACTED]  
8 [REDACTED] and [REDACTED],  
9 and that's what we effectively upload into OFA and  
10 then eventually to DYI.

11 I don't know if that answers your question  
12 or what you're looking for.

13 MR. GARRIE: So you store it in the Hive  
14 and then you use Dataswarm to schedule and that  
15 runs the processes that then pipes the data out  
16 into --

17 MR. PATEL: It has to be -- it has to be  
18 [REDACTED] Because you could  
19 think of it the data is stored at that level. [REDACTED]  
20 [REDACTED]  
21 [REDACTED]  
22 [REDACTED]

23 MR. GARRIE: [REDACTED]  
24 [REDACTED]

25 And that's done in Hive and then piped out

1 to -- where is that piped to?

2 MR. PATEL: That's then [REDACTED]

3 [REDACTED]

4 [REDACTED]

5 [REDACTED] That is, since all the data is  
6 now on a [REDACTED]

7 [REDACTED]

8 MR. GARRIE: And the data is stored in --  
9 (inaudible).

10 THE COURT REPORTER: We can't hear you.

11 MR. PATEL: Sorry, Special Master Garrie,  
12 we lost you.

13 MR. GARRIE: Sorry. I was just thinking.

14 So then [REDACTED]

15 [REDACTED]

16 [REDACTED] I guess?

17 MR. PATEL: Yeah. Correct. Like, you can  
18 think of the data, when you store the data in Hive,  
19 you're storing it on commodity kind of hardware, so  
20 like think of hard disks; and then when you're  
21 storing it in [REDACTED] you're storing it -- and I'm  
22 going to speculate here. I don't know the exact  
23 internals of [REDACTED] but I assume you're storing it

24 [REDACTED]

25 [REDACTED] [REDACTED] [REDACTED] I'm

1 not exactly sure, but it's not a hard drive.

2 MR. GARRIE: So then -- so then is there  
3 then documentation about how Facebook processes the  
4 data they get from business tools or is it by  
5 product?

6 MR. PATEL: It is by product.

7 MR. GARRIE: There is documentation for the  
8 conver- -- off-line conversion business tool that  
9 explains what you just said effectively?

10 MR. PATEL: What I explained is how we take  
11 and summarize the data in a manner that can be  
12 produced back through DYI and the OFA tool. There's  
13 a lengthy process that goes into matching users.  
14 That's not part of, like, my expertise there.

15 MR. GARRIE: Okay.

16 MR. PATEL: You are effectively taking the  
17 output of what the team is -- the product team has  
18 done and producing a summary of it.

19 MS. RING: If I may, Special Master Garrie,  
20 just again going to kind of what Mr. Patel focuses  
21 on and I don't know if this is clear, the  
22 distinction between.

23 Do you know what the OFA tool is? I don't  
24 know if that's ever been discussed with you. That's  
25 what Mr. Patel is referring to. It's another

1 transparency tool that contains a subset of OFA  
2 information. So everything is -- you know, the OFA  
3 information that goes into DYI, the OFA tool has  
4 like a subset of that data, but I just want it to be  
5 clear.

6 MR. GARRIE: Yeah. It's an off-line -- I  
7 mean, I thought that was the acronym for the  
8 off-line activity.

9 MS. RING: It is. It is.

10 MR. GARRIE: It's different than the  
11 off-line conversion tool we're talking about.

12 MS. RING: Well, the off-line tool.  
13 Off-line conversions are one thing, then there's  
14 off-line tool.

15 MR. GARRIE: I know. That's what I was  
16 talking about. Yeah. I get it.

17 MS. RING: Okay.

18 MR. GARRIE: But I'm talking about off-line  
19 tool. I'm talking specifically about the business  
20 tool, off-line tool that was identified in the  
21 answer.

22 MR. PATEL: Sorry. Maybe I'm not getting  
23 this straight.

24 Like, the off-line conversions tool is just  
25 an API that advertisers can call into.

1 MR. GARRIE: Okay.

2 MR. CLARK: Let's step back for one second  
3 and we'll walk through this. That way we'll make  
4 sure we're all on the same page.

5 So one of the business products that  
6 somebody can integrate with is actually called  
7 off-line conversions. Off-line conversions means I  
8 either have an event or a store and I provide sets  
9 or details or create a custom audience from that  
10 off-line event to drive that. We have documentation  
11 on that process and how to do that on the business  
12 tools page.

13 In the same way, we have on-line  
14 conversions that can be measured or that can be  
15 implemented via SDK, whether -- I mean, there's  
16 different types of calls and things associated with  
17 both visiting that advertiser and certain kinds of  
18 events.

19 From a transparency perspective and going  
20 to the other side for what we provide to the  
21 consumer, we have a product called OFA.

22 MR. GARRIE: That part I get.

23 MR. CLARK: Okay. And subsets of -- and  
24 sorry.

25 All of OFA is included in DYI. Some of OFA

1 is included in AYI. I just wanted to make sure we  
2 were all talking about --

3 MR. GARRIE: No. No. That was the  
4 documentation. That was the documentation you gave  
5 me. Yeah.

6 I was specifically interested in the  
7 off-line conversion tool because you referenced it  
8 and I was just wondering, but it's an API  
9 basically.

10 What I wanted to know isn't about the OFA  
11 tool. My question, sorry, I should have been clear,  
12 off-line conversion business tool is -- what I was  
13 trying to figure out is how does -- it collects  
14 data, right, from off-platform?

15 MR. PATEL: Yeah. Just one clarification.

16 Like, Facebook receives the data from the  
17 advertiser. Right. Like, there needs to be an  
18 action of a coral command run or an upload of an  
19 Excel sheet. Like that's -- that's a distinction.

20 MR. GARRIE: That's what I thought.

21 So if you have to upload it, do you know  
22 where -- where is it stored? Like, that's what I'm  
23 wondering is where is that information stored? Is  
24 it stored in the Hive or is it stored somewhere  
25 else?



1 MR. PATEL: Yeah. So it's on a per-product  
2 basis, but eventually it is stored in Hive, which is  
3 a source for the OFA tool is what I explained.

4 MR. GARRIE: I get it. I get it. Yeah.  
5 So I get that part.

6 I'm interested in the off-line conversion  
7 tool itself, not the OFA part, because the off-line  
8 conversion tool is collecting data. Basically I  
9 understand you collect data. I upload. I run a  
10 business. I upload -- put all these widgets to  
11 whoever. I upload a spreadsheet. You map it to  
12 whatever digital ad units, whatever. That then gets  
13 displayed back to me as the business owner, so to  
14 speak, to have my strategic ad spend.

15 I'm not sure exactly how it works, but some  
16 ROI that makes you want to buy more ads effectively  
17 in some way or add some ROI to the process of the  
18 money I'm spending and I'm just wondering where does  
19 that data actually get stored that they're uploading  
20 for that tool?

21 MR. PATEL: Yeah. So it would be a variety  
22 of systems. Like, our return on ad spend, for  
23 example, is probably stored in some of our [REDACTED]

[REDACTED]

[REDACTED]

1 MR. GARRIE: No. I get it.

2 So the off-line conversion tool, when they  
3 upload their data into the system, what I'm trying  
4 to understand is what data is being -- so that's  
5 data being collected via a business tool, right, as  
6 well as the ad.

7 MR. PATEL: Uh-huh.

8 MR. GARRIE: But that data being collected  
9 is data that the user's uploading.

10 So my question is, is the data collected --  
11 is that -- that particular off-line conversion tool,  
12 is that data being collected and processed, like  
13 extracted and processed on all of the data pipelines  
14 via the [REDACTED] is my question?

15 I was just trying to figure out how that  
16 tool ran. If you don't know, it's completely  
17 acceptable.

18 MR. PATEL: No. I wouldn't know the  
19 internal workings of it.

20 MR. GARRIE: That's fine. You're talking  
21 about OFA and off-line conversion.

22 All right. Next. So then broadly speaking  
23 in the answers, I'm moving forward, how many  
24 advertisers provide content lists for generating  
25 custom audiences? Like, do you track that or how

1 that comes through? I couldn't figure that out.

2 MR. CLARK: Let me see if I can clarify  
3 this because I think -- and I'll go back to that  
4 lens of like as a consumer when I download DYI, I  
5 can see every custom audience that my information  
6 was attached to or who were the advertisers that  
7 did that. We don't actually see your information.  
8 We match your information because we receive it  
9 hashed.

10 MR. GARRIE: So then is there documentation  
11 on how Facebook processes the data they get from the  
12 advertiser?

13 MR. CLARK: There is -- so I'll answer this  
14 in two ways.

15 There is clear documentation on the  
16 business tools page on how to upload and build your  
17 custom audience and how to work with that part of  
18 the process. And so that -- that documentation does  
19 exist, how to set all that up and what fields are  
20 accepted.

21 I believe the other part that you're asking  
22 for is, is there a document that explains how that  
23 works internally and where it's stored. This is  
24 going to be another -- like, it's product by  
25 product. So we'd need to get a specific engineer in

1       that space to talk about that.

2               MR. GARRIE: I got it.

3               So my question -- so to give you an  
4       example. I'm a user. I get uploaded. That gets  
5       distributed to Facebook. So it sounds like that  
6       product specifically has made -- is input into a set  
7       of schemas that is product specific that then  
8       aggregates itself out basically.

9               MR. CLARK: Correct. Especially since it's  
10      not access -- we don't -- you know, my email  
11      mike@mike.com is not going to be what Facebook  
12      receives.

13              We will receive a hash. And so that  
14      matching, like, occurs. So that one especially is  
15      processed in its own unique way and we'd want to  
16      talk to the product engineer from that space.

17              MR. GARRIE: And do we know where the  
18      storage location for custom audiences is? That was  
19      the other thing I couldn't -- I know it ends up in  
20      the DYI file at some point, but where is it actually  
21      stored?

22              MR. CLARK: I don't have that answer off  
23      the top of my head. I would need to clarify -- I  
24      would want to clarify that, unless you do, Mayur.

25              MR. PATEL: Yeah. So custom audiences are

1 stored in [REDACTED] It's just a different type of

2 [REDACTED]

3 MR. GARRIE: And are the custom audiences  
4 then stored per advertiser or how are they stored?

5 MR. PATEL: That I wouldn't know. Like, I  
6 wouldn't know the schemas of how these audiences  
7 are. I just have a general idea of, like, that  
8 system stores them.

9 MR. GARRIE: But are all the -- and the  
10 reason why I'm asking, what I'm trying to figure  
11 out, and maybe I should have asked this first, are  
12 all custom audiences included in the DYI tool in  
13 total, all of them?

14 MR. CLARK: That you are a part of or  
15 associated with. I think specifically in the tool  
16 it's advertisers using your activity or information  
17 and, to the best of my knowledge, that is the  
18 complete set of custom audiences that you are  
19 matched to.

20 MR. GARRIE: I just want to make sure I  
21 understood it.

22 If I'm a user and I end up in 20 different  
23 custom user -- custom audiences, every single one of  
24 them is included in the DYI file?

25 MR. CLARK: Correct. To the best of my

1 knowledge, yes. And we classify it as a list  
2 uploaded or used by the advertiser in the DYI file.

3 MR. PATEL: Yeah. There's a second  
4 category, too. Website or a custom -- website or  
5 app event activity.

6 MR. CLARK: That is correct. Interactions  
7 you may have had with the advertiser's website, app  
8 or store, which would include your off-line  
9 information that you're referring to.

10 MR. GARRIE: Yeah. Exactly.

11 So then you get off-line information from a  
12 third-party integrated product about whatever.  
13 That's stored by product somewhere in the data  
14 warehouse, warehouses. And then from there, whether  
15 it's from a third party or on platform, I get -- if  
16 I get tagged in a custom audience, all of that's in  
17 the DYI file.

18 MR. CLARK: Correct.

19 MR. GARRIE: All right. That was  
20 Category 2, off-platform.

21 Do plaintiffs have any questions they wish  
22 or further clarifications they think would be? I  
23 see no hands, so I assume no.

24 Okay. So then could you write a query, out  
25 of curiosity? Well, I guess --

1           So then inferred data, how is inferred data  
2 actually computed? Because I could not figure that  
3 out from the answer.

4           MR. CLARK: So it's derived and it's  
5 derived depending on the type of things. I think  
6 like there's -- I'll use a couple concrete examples  
7 because I think it will make it easier to  
8 understand.

9           I think one is your primary location. We  
10 don't ask you your primary location, but based on  
11 your location, we can derive that that is your  
12 default or primary location.

13           Your time zone, we don't ever ask you  
14 specifically for your time zone, but we're able to  
15 derive that based on your primary location and based  
16 on other sets of information.

17           And interests is one that's in the list as  
18 well that based on -- based on information, we're  
19 able to derive a set that matches to that. Even  
20 language, like watching -- or being able to derive  
21 what your preferred language is or language you may  
22 know based on responding in comments using other  
23 languages, able to derive a set of information like  
24 that.

25           MR. GARRIE: Is there a list of the data

1 sources used that identify the derived datasets?  
2 Because derived data varies by definition, but it's  
3 another data; right? So is there --

4 MR. PATEL: Is there a particular type of  
5 data that you're interested in?

6 MR. GARRIE: Well, for all of us, the judge  
7 said all data on derived. It wasn't limited.  
8 Derived data was very broadly defined. So I would  
9 welcome any insights, whether narrow or broad. The  
10 judge did not delineate derived data from geo  
11 location or from phone or from third party. It  
12 wasn't delineated.

13 MR. PATEL: Our ads interest, I can speak  
14 to, are derived from first-party activity. So any  
15 activity that you perform on Facebook. So liking a  
16 page, if you like swimming pages or if you like  
17 Rebecca Soni, which I do, then you -- you know,  
18 those are ads interest of swimming that you can be  
19 put into. Or, you know, clicking on specific ads,  
20 that's information that we derive from your  
21 activity.

22 The location depends on IP address,  
23 potentially the hometown you put, you know, state  
24 in your profile and other information could be  
25 pulled in. I'm not the expert in that. Locale is



1 probably from your device settings or your browser  
2 settings and which language you pick and we want to  
3 make sure that, you know, we produce the right  
4 communication mechanism back to the user there as  
5 well, the right language.

6 So these are just broad set of examples.

7 MR. GARRIE: So I guess what I'm trying to  
8 figure out is, what is the source of the data that  
9 is derived data? It's all -- like, is it -- I mean,  
10 I get it's computed based on some algorithmic  
11 analysis that you perform.

12 MR. CLARK: It connects back to one. Like,  
13 my on-platform activity, that is the mix of --  
14 that's both -- from a category perspective, that's  
15 both the user-provided data, the things like  
16 preferred location, like the first place is what do  
17 you put as where you live right now, but also  
18 observed data and that observed data is IP address,  
19 it is everything else that I interact with.

20 So I believe -- I believe mine currently  
21 says that I live in -- my location as it says in my  
22 DYI is Denver, which is accurate, but I believe my  
23 profile says I live in Oakland. So it's based on  
24 this combination of first-party activity as we  
25 outlined in one for this derived data.

1 MR. GARRIE: What about off-platform  
2 activity as well? Does the off-platform activity  
3 factor into derived data analysis?

4 MR. CLARK: Not --

5 MR. GARRIE: Because you put examples in  
6 there about like music or recommendations.

7 MR. CLARK: A couple of things that we  
8 refer to as -- and this is where, like, defining on  
9 or off platform, things that aren't Facebook, like  
10 the portal, the frames that we talked about, my  
11 voice interactions are derived data that I have  
12 access to. It's a product that we make, but it's  
13 not on the Facebook platform.

14 MR. GARRIE: So I think I know where this  
15 goes. I think I've finally figured out what Mayur  
16 was telling me before.

17 But once a data point about me is derived,  
18 where does it usually end up? I assume it's either  
19 the TAO or Hive, but I don't know that.

20 MR. ZARASHAW: So city ends up being  
21 platform specific and usually depends on how the  
22 data will in the future be used. Usually the  
23 decision is based on do we need to use the data in  
24 real time in the product to inform some other part  
25 of the user experience, in which case it goes into

1 past expansive storage. So, for example, TAO. If  
2 it's something that will only be used for off-line  
3 analysis, it goes into slow cheap storage, such as  
4 Hive.

5 MR. GARRIE: I mean, would it be a fair --  
6 is there a list of all -- is there a list of all the  
7 derived data points being calculated for a user?

8 MR. ZARASHAW: It, once again, ends up  
9 being something very product specific. I'll build  
10 on Mike's earlier new user experience example.

11 You can imagine that you can do a big  
12 pop-up experience the first time a user visits  
13 Marketplace to walk them through every aspect of  
14 how Marketplace works and it's going to be an  
15 intrusive experience because you are keeping the  
16 user from using the product until they go through  
17 the tour and it would be rather annoying for a  
18 person if that popped up every single time they went  
19 to Marketplace after the first.

20 So one piece of derived data would be has  
21 the person gone through that tool before. Whoever  
22 is implementing it on the Marketplace team would  
23 decide how should they store that so that every  
24 single time the person comes to Marketplace we check  
25 it; and if the person has already gone through the

1 tour, we don't pop it every single time.

2 MR. GARRIE: That would be stored in TAO;  
3 right?

4 MR. ZARASHAW: Exactly. So, most likely,  
5 TAO because we would not be able to retrieve it fast  
6 enough from Hive.

7 MR. GARRIE: Yeah. So to repeat my earlier  
8 statement, I'm not looking for 100 percent accuracy  
9 here. I'm just looking for the systems that, most  
10 likely, have -- like, it sounds like every piece of  
11 data you have would -- is a draconian process for  
12 an individual person, but it sounds like it either  
13 ends up in TAO or some subset of tables in  
14 the (inaudible).

15 THE COURT REPORTER: I'm sorry.  
16 Mr. Garrie, we can't hear you.

17 MR. GARRIE: Okay. Yes.

18 So depending on the utility and value of  
19 the data, it, more or less, will sit in TAO or  
20 Hive.

21 MR. ZARASHAW: I wouldn't say utility or  
22 value, but just how it would need to be accessed in  
23 the future. Yes, most of the time it will end up in  
24 TAO, Hive or both.

25 MR. GARRIE: And you would need it

1 sometimes in both because you're doing integrity  
2 analysis in Hive and you would need that data and  
3 you wouldn't want to take the computing cycles and  
4 hit.

5 MR. ZARASHAW: Precisely.

6 MR. GARRIE: Does Plaintiff have any  
7 further questions about Category 3, inferred data?

8 MS. WEAVER: I think that we do, but I  
9 think we need to caucus and we'll circle back. But  
10 you should move on, Special Master Garrie.

11 MR. GARRIE: I'm going to keep going. When  
12 we take our next break, that will be an opportunity  
13 to caucus and I'll consider the questions at that  
14 point.

15 MS. WEAVER: Thank you.

16 MR. GARRIE: Can we keep going? Is that  
17 okay for everybody? It's good for me. Is it good  
18 for everyone else? Say if it's not, then we'll go  
19 from there.

20 MS. RING: Yes.

21 MR. GARRIE: Perfect.

22 Now we're getting to, so there's this  
23 paragraph and I don't know, it's a question I had,  
24 which is what are the reasons for not including  
25 content specific data in the DYI tool?

1 MR. MITCHELL: Could you clarify what you  
2 mean by content specific data?

3 MR. GARRIE: I could read you as it was  
4 described to me in the answer and then I will.

5 It says -- I just want to be -- I want to  
6 be --

7 Like what item I added to the cart.  
8 Remember that earlier example about I'm shopping and  
9 I add that to the cart. That doesn't end up in the  
10 DYI file, but it's shared back to you guys.

11 MR. CLARK: So perhaps the place to start  
12 is looking at what is in the DYI file.

13 MR. GARRIE: I did that actually. So I got  
14 that down.

15 MR. CLARK: Okay. So when you look, and  
16 I'm just going to look at my own personal one,  
17 the -- I'll go to Luluemon and I will see that I  
18 have page click. I will also see that I have, well,  
19 page view, initiate checkout, and purchase as event  
20 names, I have the date and time that those were  
21 received on, and then I have an identifier.

22 And going back to the documentation on the  
23 business tools that we have, there are primary  
24 things that we map that are translatable that we can  
25 interpret based on how the information is given to

1 us and that information are the fields. We know  
2 that this -- like, one, it's tied to myself as a  
3 user. There is an IP to this event. That event  
4 included the initiate checkout and received on.

5 We do have custom fields, but we don't have  
6 control over what the person puts in those custom  
7 fields nor are they translatable to us. If you've  
8 gone and looked at examples yourself, they don't map  
9 to specific schema.

10 MR. GARRIE: Yeah. But it does not include  
11 data identifying the specific content that was  
12 viewed for the item that was added to a cart.

13 So my question is, what are the reasons for  
14 not including --

15 You asked me specific content. So I'm  
16 giving you examples.

17 So what are reasons for not including  
18 content specific data in the DYI tool?

19 MR. PATEL: It's not interpretable by us.  
20 It could be any random string that the advertisers  
21 could send. Also, the events that are in place,  
22 like the view content, the add to cart is something  
23 that the developer builds. Right. Like, not all  
24 developers have the same competency as Luluemon.  
25 There could be developers that accidentally just

1 say you purchased everything instead of viewing  
2 content. We can't verify all the inputs there, and  
3 so it's really, like, unintelligible to us to  
4 explain that to users.

5 MR. GARRIE: Okay. So that I understand.

6 MR. PATEL: Okay.

7 MR. GARRIE: But you do -- but the company  
8 provides you that data and I get that, the integrity  
9 of the data, the quality of the developer, and the  
10 other factors you've explained. My question is more  
11 different.

12 So that data is stored somewhere; right?  
13 And so what's the decision -- how is the decision  
14 made about the content specific data provided by  
15 like Luluemon is a great example that there was a  
16 decision made that, okay, this could go in the DYI  
17 file, but the other person, who knows, whatever,  
18 didn't make it in? So like content specific.

19 MR. CLARK: We don't filter out by -- like,  
20 Costco and United and Luluemon are all there and  
21 we're not filtering out anybody that sent events if  
22 that's -- I want to make sure, Special Master  
23 Garrie.

24 MR. GARRIE: No. I'm saying content  
25 specific. So I'm not saying you're filtering out



1 by vendor.

2 Is there any content specific things that  
3 aren't included in the DYI tool that you get that  
4 isn't included?

5 MR. PATEL: Yeah. So there would be the  
6 event names if you have custom event names are the  
7 content. We have standard names like view content,  
8 purchase, add to cart, leave, and there's a few  
9 others. I can't name them all off the top of my  
10 head.

11 MR. GARRIE: But they're documented.

12 MR. PATEL: They're documented in the  
13 public API, the public documentation. But an  
14 advertiser, if you think of it, a Pixel Fire or an  
15 app event is just a coral call. Like, it's just an  
16 HTTP request that goes over and can be modified to  
17 anybody as well. Right. Like, you can add a  
18 custom -- like, is this a JSON log or maybe it's in  
19 the query params, but you can add any kind of event  
20 name that you want. It's not -- if you choose, you  
21 want to call it add to cart without the underscores,  
22 that's up to you. But we provide a standard set of  
23 features. And if we get something that's not in the  
24 standard for the event name, we call it custom  
25 events because we just don't know what's in there.

1           And similar for content we provide. Yes,  
2     you come and say what the value is, what the  
3     currency is as well as what -- what content you may  
4     have added as well as the name, but we can't verify  
5     that at that point.

6           MR. GARRIE: That doesn't end up in the DYI  
7     file.

8           MR. PATEL: Correct.

9           MR. GARRIE: So I'm just thinking about it.  
10    It's just a lot of information and I just want to  
11    make sure I understood.

12           So then are there any reasons besides it  
13    not being intelligible, or it's not it's not  
14    intelligible, they didn't tell you, not including  
15    content specific data in the DYI tool?

16           MR. PATEL: No. Besides that, no.

17           MS. RING: Special Master Garrie, I just  
18    want to clarify.

19           When you said not intelligible, they tell  
20    you, what they're saying is what they tell us is not  
21    intelligible. So I just wanted -- I want to clarify  
22    if you're thinking these are two different things  
23    or --

24           MR. GARRIE: They didn't tell --

25           MS. RING: -- let them explain something.

1 MR. GARRIE: What happened is they didn't  
2 bother to tell Facebook what they're doing, and so  
3 Facebook doesn't know how to interpret the custom  
4 values and, therefore, it doesn't stick them in the  
5 DYI file. So it's unintelligible to Facebook  
6 because nobody's told Facebook what the data is that  
7 they're getting from the third party because they  
8 have no way to verify because they have a standard  
9 set of APIs that are public that say this is how it  
10 works, but they have custom events or JSON. I don't  
11 think it's JSON, but they have custom events that  
12 could be called and that information, they don't  
13 know what it is, isn't included in the DYI file, but  
14 Facebook gets the data.

15 MS. RING: But the data is gibberish  
16 sometimes. I just want to make sure that  
17 distinction is not lost.

18 MR. GARRIE: Gibberish means that it's not  
19 intell- -- to the person that's providing it, it  
20 makes sense. To Facebook, it doesn't make sense  
21 because nobody told Facebook what it is.

22 MS. RING: I defer to the engineers on  
23 that, but I think what we saw is gibberish. So I  
24 don't know. How would you guys explain it? Sorry.

25 MR. CLARK: I would assume it has meaning

1 to the person that wrote that integration.

2 MR. GARRIE: Yeah. I mean, I would assume  
3 that's why they wrote it.

4 I mean, the real question, I guess -- I  
5 mean, maybe there are developers that just  
6 write gar- -- yeah, it's probably true there are  
7 developers that just write garbage for writing code,  
8 but assuming there are developers that --

9 I mean, the real question is, you store all  
10 custom event data that is -- I mean, or do you  
11 discard it is the question, I guess, for the user?

12 MR. CLARK: I would want to -- unless Mayur  
13 has a specific answer, I would actually want to talk  
14 to -- I'd want to -- like from this -- this is the  
15 extent of my understanding on that process and  
16 product and how it worked on it. I'd want to talk  
17 to that product team specifically.

18 MR. GARRIE: Unless we can get a quick  
19 answer. No. I'm kidding. No pressure.

20 So we'll talk to the product team and we'll  
21 go from there. All right. That makes sense to me.

22 All right. So then question number three,  
23 I promise not all of them will -- they will speed  
24 up. But at the beginning and hopefully we don't  
25 have to do it again. So better now than do it

1 twice.

2 What are some examples of the Dataswarm  
3 tasks? I saw about 10 of them. If you want, we can  
4 go through. Just at a high level, what are some  
5 examples, and then we can look at --

6 MS. RING: Special Master Garrie, none of  
7 these engineers here are specialists or, you know,  
8 deal with the Dataswarm. We were just answering  
9 this question, and so we put together -- we were  
10 trying to answer the questions. We put together --  
11 we first tried to identify what the system was, and  
12 then we gave you the samples just so we can show you  
13 what's there, but none of the engineers that are  
14 here today, since that wasn't one of the categories  
15 for engineers, focuses on Dataswarm.

16 MR. GARRIE: Okay. I'll just tee up the  
17 questions I do have so you can --

18 MS. RING: Okay.

19 MR. GARRIE: -- at least figure out.

20 MS. RING: Okay.

21 MR. GARRIE: So, you know, does a task  
22 operate on a single data point or on multiple rows?  
23 That's one question I'm interested in.

24 Is there documentation on task dependency  
25 chains that were created for specific purposes?

1 Based on what I've heard, I'm pretty certain there  
2 are because it sounds sequential at least for some  
3 of it, but I could be wrong.

4 And then under that is how does Facebook  
5 keep track of what each task is used for? Because  
6 it is -- you know, and are tasks ever modified?

7 And then are dependency -- this is  
8 something that -- are dependency chains ever  
9 removed to reduce the amount of computing that is  
10 done?

11 And then how does Facebook keep track of  
12 what all the Dataswarm tasks accomplish?

13 And then is there a task -- and I don't --  
14 well, I'll ask the question.

15 Is there a task for each Hive table that  
16 serves as a start task, in quotes, like a start  
17 task? Because what I saw when I was looking at it,  
18 it seemed to suggest that it was a possibility.

19 And then do Dataswarm tasks write data to  
20 other systems or only back into the Hive? I think I  
21 know the answer based on what we've talked about,  
22 but it would be good if there was someone that  
23 could.

24 And then is the data in the Hive used by  
25 production systems? I think -- yeah.

1           And then how do production systems query  
2 data in the Hive? I think it's [REDACTED] Maybe not.  
3 I don't know. And if so, you know, I have a couple  
4 questions about [REDACTED] because this all writes  
5 directly to the -- eventually this is all user data  
6 that's being operated on in the Hive.

7           And then how is data from the Hive  
8 transferred to other storage that can be queried? I  
9 think I got some of the ways. Sounds like they can  
10 do it any way they want, but maybe I'm wrong.

11           So those are sort of the questions I had  
12 about the Hive and [REDACTED] if my assumption's right.

13           Well, that was fast. Let's go to question  
14 four.

15           So I guess, broadly speaking, before we go  
16 to question four, does Facebook retain user activity  
17 that doesn't end up in the DYI file?

18           MR. CLARK: As I walked you before, like  
19 there are specific cases of data where like that  
20 association isn't -- like we can't fully determine  
21 or guarantee that it's that user and wouldn't want  
22 to give them somebody else's data.

23           MR. GARRIE: So like a constant event  
24 basically.

25           MR. CLARK: I don't know that. That may

1 not be the example, but if that helps a mental model  
2 to think about it.

3 MR. GARRIE: Okay. So there's basically  
4 the same answers you told me in the beginning. I  
5 get it.

6 MR. CLARK: Yes.

7 MR. GARRIE: Question four, the user ID. I  
8 guess, does the information -- I don't know -- I  
9 don't know if this is the right person, but is the  
10 information about the user, first name, last name,  
11 email, phone number, stored in one specific  
12 location?

13 MR. PATEL: Only one specific location or  
14 do you mean a specific location that you can go look  
15 up?

16 MR. GARRIE: One or the other.

17 MS. RING: Wait. I'm sorry.

18 Just so it's clear, question four is about  
19 identifiers. So like identifiers. Not information  
20 about --

21 MR. GARRIE: But a user ID -- but a user ID  
22 is information about the user, such as first name,  
23 last name, email, phone number that would be  
24 associated with the user ID, but it could match  
25 their user table.



1 MS. RING: I just want to clarify that, you  
2 know, this is not what the question was about, but I  
3 think they were asked to answer internal identifiers  
4 that Facebook was using, just to be clear.

5 So you're asking if Facebook uses name or  
6 phone number as an internal identifier?

7 MR. GARRIE: No. Maybe I shouldn't have  
8 said name.

9 My question is, is the information about  
10 the user ID stored in one specific location?

11 MS. RING: Okay. So it's about  
12 identifiers.

13 Do you guys know?

14 MR. ZARASHAW: So user IDs are 64-bit  
15 integers where a set of databases that are back in  
16 TAO are used to generate the user IDs and they are  
17 used as a foreign key everywhere else to refer to  
18 that user.

19 There is a primary place where that user ID  
20 and some information about a user is stored, which  
21 is TAO. However, the information may also be stored  
22 in other systems and duplicated.

23 MR. GARRIE: That makes sense.

24 So my question is, is there a master user  
25 ID table, I guess? It sounds like there is.

1 There's 64-bit characters and you're running out of  
2 them. So you must have been there.

3 MR. ZARASHAW: And there's no single --  
4 it's not a single table. It's a distributed that  
5 handles the assignment and mapping of user IDs with  
6 other core functions of TAO.

7 MR. GARRIE: I forgot you're so big. Yes.  
8 I know there's a given size. Yes. Just like TAO's  
9 the extraction layer in itself.

10 So then do any other systems store this  
11 information or do they only store user ID?

12 MR. PATEL: And when you mean information,  
13 do you mean the first name, last name?

14 MR. GARRIE: It's whatever you're putting  
15 in the user ID.

16 MR. PATEL: Yeah. The data could be  
17 duplicative in a warehouse, for example, when we're  
18 doing analysis.

19 MR. GARRIE: And what other systems can  
20 access? Do you know a way of figuring out what  
21 other systems can access the user ID information?

22 MR. ZARASHAW: Special Master Garrie, what  
23 do you mean by access the user ID in this case?

24 MR. GARRIE: Well, if you use it and  
25 extract it to TAO, right, TAO is -- I guess from

1 what I understand is -- if there's a master user  
2 table and then that's extracted to like a user ID  
3 table, is it protected information; right? Is the  
4 user ID information limited access, I guess?

5 MR. CLARK: So across all of our systems --

6 MS. RING: Sorry. I just want to clarify.

7 I didn't hear anyone say a master user  
8 table. I just want to be really clear what we're  
9 talking about.

10 MR. GARRIE: User ID.

11 MS. RING: There's a table that has user  
12 IDs, but no one -- I mean, I just want to make sure  
13 we're being clear.

14 MR. GARRIE: Well, it's a 64-bit  
15 distributed.

16 MS. RING: That's for the user ID, but --

17 MR. GARRIE: Correct.

18 MS. RING: Okay.

19 MR. GARRIE: But the user ID is associated  
20 with user information. No?

21 MS. RING: I don't -- I think what we  
22 heard, there's no one here that knows exactly  
23 whether that's the case.

24 I think Eugene was just saying that the  
25 user ID is stored as a 64-bit. I can't even repeat

1 it, Eugene. I'm sorry. But I just want to make  
2 sure that we're not, you know, misunderstanding here  
3 what actually is being talked about, especially when  
4 we're talking about then what privacy safeguards are  
5 applied to it.

6 MR. CLARK: I have a caveat because, you  
7 know, part of -- part of having TAO as that  
8 extraction layer is to manage access like for any  
9 kind of data access, you know, that goes back to as  
10 engineers design systems, like they are trained in  
11 specific things that they are supposed to build,  
12 and before we can access certain kinds of data or  
13 have any kinds of things, those things have to be  
14 defined and mapped. And so that -- that mapping is  
15 to the structure, but not necessarily the data in  
16 it, and that data is accessed as needed for the  
17 functionality of the product.

18 MR. GARRIE: So then who owns the user ID?

19 MR. CLARK: It's no one and everyone.

20 MR. GARRIE: Well, who gets fired if you  
21 lose the table?

22 MR. CLARK: Well, that's part of being able  
23 to operate at this scale, we wouldn't have a single  
24 table where you --

25 MR. GARRIE: It's distributed. Right.

1 Sorry. No. No. Sorry. I keep forgetting it's  
2 distributed to like where --

3 So the user ID that you guys offer as a  
4 definition is Facebook uses an industrywide  
5 technique called pseudonymization to represent users  
6 on the Facebook platform. You create a 64-bit  
7 string.

8 Who owns that process? So there should be  
9 a single source of truth for mapping the UID to --  
10 no? Isn't there a single source?

11 MR. ZARASHAW: I'm not sure of who even  
12 owns the process. It is a piece of infrastructure  
13 that has been functioning quite well for about a  
14 decade and has not really needed significant  
15 changes.

16 MR. GARRIE: Yeah. What I'm asking is  
17 what's the single source of truth for mapping?

18 MR. ZARASHAW: TAO.

19 MR. GARRIE: TAO. But that's for every --  
20 that's your distributed table database system, but  
21 that's it?

22 MR. ZARASHAW: Yes.

23 MR. GARRIE: That's your single source of  
24 truth?

25 MR. ZARASHAW: Yes. So if you go to

1 Facebook and to sign in and we go to validate your  
2 password and so on, that will go from TAO.

3 MR. CLARK: If you're an engineer building  
4 your primary storage interface as we talked about  
5 like looking at the 95 percent case, that  
6 interaction is going to be through TAO.

7 MR. GARRIE: So moving to the replacement  
8 ID or the RID. It's very interesting.

9 So as long as the RID to user ID mapping is  
10 not deleted, can the data in the Hive be mapped to a  
11 specific user?

12 MR. PATEL: Yes. If the account is not  
13 deleted.

14 MS. RING: Okay. Sorry. Are you saying  
15 all data, Mayur, all data in Hive? Let's just be  
16 really precise especially on this topic because  
17 we've spent so much time on it previously.

18 MR. PATEL: Yeah. Any data associated to a  
19 non-deleted user if it's stored within RID can be  
20 associated back.

21 MS. RING: Right. If it's associated.  
22 Again, let's just be clear about what we're talking  
23 about.

24 MR. GARRIE: That's what I said. No,  
25 Counsel Ring.

1 MS. RING: Okay. I didn't hear that. I  
2 don't think it was phrased exactly that way. So I  
3 just wanted to --

4 MR. GARRIE: It was.

5 MS. RING: If that's what you meant and now  
6 we're on the same page, great.

7 MR. GARRIE: So let me repeat my question  
8 so there's no confusion.

9 So as long as the RID to user ID mapping is  
10 not deleted, can the data in the Hive be mapped to a  
11 specific user? I assume the answer is yes, which is  
12 why you built it that way.

13 MR. ZARASHAW: So it's complicated because  
14 it depends for a specific Hive table what is its  
15 schema. There exists Hive tables that use user IDs  
16 in its schema, there are Hive tables that use  
17 replacement IDs in its schema, and there's Hive  
18 tables that don't have any kind of user identifier,  
19 RID, SID.

20 MR. GARRIE: I get that.

21 I'm asking if RID to user ID is not  
22 deleted.

23 What I'm saying is if that mapping still  
24 exists -- I'm not talking about all the scenarios  
25 you're talking about. I'm talking about RID to user

1 ID mapping is not deleted, meaning that it exists,  
2 can the data in the Hive be mapped back to a  
3 specific user?

4 MR. ZARASHAW: I'd like to get a little  
5 more precise about which data in the Hive because  
6 even if you --

7 MR. GARRIE: Say the Hive is --

8 My question is, tables in RID columns. Not  
9 all tables.

10 MR. ZARASHAW: Okay. That makes sense.

11 As long as the mapping still exists, the  
12 user still exists, and a table contains an RID  
13 column, yes, it can be mapped.

14 MR. GARRIE: Yeah. Okay. So all right.

15 MR. CLARK: And the 90 days is the  
16 commitment for that deletion. That can happen  
17 sooner.

18 MR. GARRIE: But there's no way to verify  
19 across -- do you have processes that run across the  
20 Hive to validate, verify deletion is happening?

21 MR. CLARK: Yes.

22 MR. GARRIE: Across the entire Hive?

23 MR. CLARK: Well, no. We have for where  
24 those mappings have occurred because, once again,  
25 operating on that scale, you don't take a whole



1 storage system and do that mapping. It's where  
2 those RIDs have been tagged that that kicks off the  
3 trigger for that deletion process and then a  
4 validation process as well.

5 MR. GARRIE: Yeah, which you know about.  
6 Yeah. You have to know about it.

7 MR. CLARK: Correct.

8 MR. GARRIE: Okay. So then moving on to  
9 the SID, the separateable which, by the way, who  
10 thought of that, a separateable ID? That was --  
11 that was -- I've seen a lot of people try to -- I  
12 don't think lawyers appreciate how that's a very  
13 nice -- that's an interesting approach.

14 So is the SID tasked in by the off-platform  
15 system or is it mapped before writing data to the  
16 Hive? I couldn't tell.

17 MR. PATEL: We do not receive SIDs from  
18 third parties. So it would be the latter.

19 MR. GARRIE: Is then the UID passed in as  
20 well?

21 MR. PATEL: Passed in where? Sorry.

22 MR. GARRIE: So is UID passed in then and  
23 then you store it as the SID? Is that how that  
24 works out?

25 MR. PATEL: Yeah. We match to the UID. So

1 this is on Facebook.

2 MR. GARRIE: Yeah, on Facebook.

3 MR. PATEL: We do not get UUIDs from third  
4 parties. We match to a UUID and convert it to an  
5 SID. Yeah. The difference is it's similar to an  
6 RID. Where the RID is the lifetime of the account,  
7 the SID is basically -- is recycled when the control  
8 is executed.

9 MR. GARRIE: What do you receive then from  
10 the third parties?

11 MR. PATEL: What we receive from the third  
12 parties depends on the business tool. It ranges  
13 from cookies to device identifiers and, in your  
14 off-line conversions case, PII. Hashed PII just to  
15 be clear on that.

16 MR. GARRIE: Then is it possible to  
17 determine the list of Hive tables that contain UUID,  
18 RID or SID as a column value?

19 MR. PATEL: It is possible.

20 And your question is for off-Facebook  
21 activity specifically?

22 MR. GARRIE: Just generally.

23 MR. PATEL: Yeah. So we do have schemas  
24 and developers are annotating columns in Hive for  
25 whether or not there is a UUID present or an RID

1 present.

2 MR. CLARK: Individually per table across  
3 many, many millions of tables. And so could we look  
4 in the phone book and find Mike Clark? Yes. Could  
5 we look in a warehouse of? And so there becomes a  
6 technical feasibility because not all that storage  
7 that sits in Hive is online. And so that would --  
8 it becomes a technical feasibility challenge to do  
9 that full discovery across everything.

10 MR. GARRIE: I didn't say what anybody was  
11 going to do yet, but I get there are millions of  
12 tables.

13 MR. CLARK: No. I was just clarifying that  
14 last bit.

15 MR. GARRIE: No. I understand. I just  
16 want to know if it's -- like is it possible and  
17 whether or not it's economical or reasonable or any  
18 of that. I understand that there are millions of  
19 tables and et cetera.

20 Counsel Weaver, do you have a question?

21 MS. WEAVER: Yes. And maybe this is  
22 further in the outline.

23 But with regard to hashed PII and  
24 conversion and cookies received from third parties,  
25 it might be helpful to know what cookies, if they

1 are Facebook cookies and if the cookies contain  
2 identifiers that Facebook can decode and any  
3 identifiers or cookies that are used for mapping or  
4 matching purposes.

5 MS. RING: Sorry. Counsel Weaver, I didn't  
6 hear the first part.

7 Cache did you say?

8 MS. WEAVER: He said hashed PII.

9 MS. RING: Hashed. Thank you. Okay.

10 MS. WEAVER: Like corned beef.

11 MS. RING: Got it. Okay. Got it.

12 I don't know if any -- you know, the people  
13 that we have here today focus on those things. I  
14 don't think so. So...

15 MR. PATEL: Yeah. I'm not a measurement  
16 tabs engineer.

17 MR. CLARK: I wouldn't be able to speak to  
18 that either.

19 MR. ZARASHAW: I wouldn't either.

20 MR. GARRIE: You've got to -- you've got to  
21 weigh in here, Ben. You got to turn mute off.

22 MR. MITCHELL: Sorry.

23 I can't speak to that either. I didn't  
24 realize that was actually an explicit question.  
25 Sorry. No.

1 MR. GARRIE: Your counsel is making a  
2 point. So I just wanted everybody to answer.  
3 Okay.

4 THE COURT REPORTER: Mr. Garrie, whenever  
5 it's convenient, can we take a break?

6 MS. WEAVER: Of course. We'll take a  
7 10-minute break. We'll go off the record.

8 (Recess taken.)

9 (Off the record at 1:09 p.m. Back on the  
10 record at 1:31 p.m.)

11 MR. GARRIE: Back on the record.

12 So question five. Starting with TAO, my  
13 first question is, is there documentation of TAO  
14 objects and associations and their schemas?

15 MR. ZARASHAW: In this case, TAO is  
16 self-documenting like a lot of the code in that we  
17 look at the definitions of the TAO schemas inside  
18 TAO itself as the documentation. The challenge is,  
19 otherwise, documentation tends to drift and not get  
20 updated in time. So the systems get updated,  
21 schemas and TAO will get updated; but if we have  
22 separate documentation, it would never be up to  
23 date.

24 MR. GARRIE: But there is documentation of  
25 the objects and associations and their schemas.

1 They're not embedded there.

2 MR. ZARASHAW: It's not really  
3 documentation. The documentation is pretty much  
4 dumping out the schema from specific tables in TAO.

5 MR. GARRIE: Printed out.

6 MR. ZARASHAW: Exactly. If you know which  
7 table you want, because there are quite a few  
8 tables, you can go and get the schema for that  
9 table.

10 MR. GARRIE: Is there a key?

11 MR. ZARASHAW: [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

17 MR. GARRIE: How many tables are there in  
18 TAO?

19 MR. ZARASHAW: I do not know the answer.  
20 Anybody else?

21 MR. PATEL: I do not know offhand.

22 MR. CLARK: I do not.

23 MR. GARRIE: So then --

24 MR. ZARASHAW: It's one of those things  
25 where every product team -- very similar to Hive,

1 any product team can just go and create one.

2 MR. GARRIE: So then how many systems --  
3 well, this comes to my next question, which is how  
4 many systems -- how many systems write data to  
5 individual objects and/or associations? Like how  
6 many systems create posts or so on and so forth?

7 MR. ZARASHAW: I would -- this would be a  
8 complete guess, but I would guess nearly all at  
9 Facebook. It's the primary method for storage of  
10 data that is then going to be used in real time by  
11 something else.

12 MR. GARRIE: So is there a comprehensive  
13 list of the systems that write to TAO?

14 MR. ZARASHAW: There is not.

15 MR. GARRIE: And there's no way to figure  
16 out a list of systems and what data they write to  
17 TAO specifically around user data or any data? I  
18 mean, we can start broadly or we can get specifics.

19 MR. ZARASHAW: It would be an incredibly  
20 large undertaking, on a similar scale to what it  
21 took to be able to download all your information.

22 MR. GARRIE: But, I guess, do the systems  
23 have their own storage as well in using TAO, I  
24 guess?

25 MR. ZARASHAW: To my knowledge, there's no

1 way to say, hey, this specific portion of TAO is  
2 related to this specific product, if that's what  
3 you're asking. [REDACTED],  
4 much like Hive, [REDACTED]  
5 [REDACTED], somebody  
6 can create a table under that name with any schema  
7 they choose and use it.

8 Effectively, it's a way of giving you  
9 distributed database like storage, graph database  
10 like storage that automatically scales to a very  
11 high level without the product engineers having to  
12 think about the scaling challenges at all.

13 MR. GARRIE: Is the data in the DYI file a  
14 subset of data in TAO?

15 MR. ZARASHAW: Mike? Mayur?

16 MR. CLARK: I'm trying to get that.

17 [REDACTED]  
18 [REDACTED] for -- as you deem  
19 more to the example of this is where [REDACTED]

20 [REDACTED]  
21 [REDACTED]

22 MR. ZARASHAW: Blob storage; right?

23 MR. CLARK: That is correct. [REDACTED]

24 MR. GARRIE: Okay. So [REDACTED]

25 And then did you want to add something,



1 Mayur?

2 MR. PATEL: Oh, all right. So I guess for  
3 [REDACTED] it's just a storage system that lives  
4 under [REDACTED] Right. It's really -- you know, the IDs  
5 that you have for that photo are accessed through  
6 [REDACTED] and eventually pulled from [REDACTED]

7 MR. GARRIE: And is TAO cued to UID?

8 MR. ZARASHAW: [REDACTED]  
[REDACTED]  
[REDACTED]  
[REDACTED]  
[REDACTED]  
[REDACTED]

12 MR. GARRIE: So then can you describe  
13 Ent schema and Node? I didn't quite understand how  
14 that's operating.

15 MR. PATEL: You're on mute, I think.

16 MR. CLARK: That will definitely not get  
17 reported if I talk on mute.

18 Ent and Node are software frameworks and  
19 they're software frameworks that are specifically  
20 built to help identify what kind of data can be  
21 accessed. Ent and Node depend on schematization of  
22 a data system in order to know how to operate and  
23 interact with it. The instruction per table --

24 MR. GARRIE: How do you --

25 MR. CLARK: Sorry.

1 MR. GARRIE: Maybe a different way then.

2 How does someone use Ent schema if I'm an  
3 engineer, concrete example, or Node?

4 MR. CLARK: So I'll finish explaining what  
5 they are because that will help you explain how an  
6 engineer might use them.

7 So, for instance, details about what can be  
8 accessed are included in that schema and Ent and  
9 Node as software frameworks are what enforce that.  
10 Ent is used specifically for the hack programming  
11 language and Node is for all the other programming  
12 language that we have at Facebook.

13 And so I, as a developer, when I go to  
14 write something, if I'm using one of the existing  
15 software frameworks, they automatically integrate  
16 with Ent and Node; but also as part of privacy  
17 programming training, if I'm writing from outside  
18 of one of the existing frameworks, I'm required to  
19 use Ent or Node in accessing the data and it makes  
20 that determination.

21 Eugene or Mayur, anything else you'd add?

22 MR. PATEL: Yeah. There is a set of  
23 privacy rules that you also set, like who can view a  
24 piece of content as well as the deletion rules on  
25 how a piece of content is deleted.

1           So, for example, if you have a comment on a  
2 photo, if you delete the comment, you shouldn't  
3 delete the photo, but you should delete any comments  
4 that were below the comments or like, I guess if you  
5 look at it in a tree, comments that were for that  
6 comment, not comments that were associated to the  
7 photo specifically, and those rules are then defined  
8 in the schema in that Ent schema Node.

9           MR. GARRIE: What would be a tangible  
10 engineering example of how they use it?

11           MR. CLARK: It's built in so that they  
12 don't have to think about how to apply the privacy  
13 controls and other elements. I don't have to  
14 understand all of the underlying storage or  
15 everything that's there. It takes care of that for  
16 me. So I don't have to specific call those. It's  
17 done with every call, if that makes sense,  
18 Special Master Garrie.

19           MR. GARRIE: Yeah, but I need like  
20 real-world examples of how it's actually used  
21 because I don't understand conceptually, like, how  
22 you use it as a developer to manage user data. Like  
23 I don't quite understand.

24           MR. PATEL: An example of this is if you  
25 created a schema for photos that people upload, you

1 would be required to implement the can see function.  
2 And I'm not sure if that's exactly the function, but  
3 like can see, we can just assume it's whether or not  
4 a user looking at the piece of content can actually  
5 view the content.

6 So an example of how you would write it is  
7 the author of the content, the photo, so if user ID  
8 is equal to the person who created it, great, show  
9 on the piece of content. And then you would go down  
10 the privacy settings of the photo. Are we allowed  
11 to show it to your friends? Are we allowed to show  
12 it to the public? Are we allowed to show it to your  
13 friends of friends? And that would have to be  
14 implemented in that can see function, and so it's a  
15 requirement that we enforce these calls.

16 MR. GARRIE: And you enforce it using --  
17 like so then -- so let's say that user ID equals  
18 this person, show them this content, and then so  
19 where does Ent schema and Node come into this?

20 MR. PATEL: Yeah. You could think of it  
21 like if somebody were accessing the data directly  
22 from [REDACTED] that would be a problem. We wouldn't  
23 check the privacy rules, obviously. So to make it  
24 scaleable, you have to go through Ent schema. Like  
25 that's a requirement.

1 MR. GARRIE: Ent schema is a way to access  
2 data. And so getting past all the privacy,  
3 whatever, Ent schema is a way to access the photo?

4 MR. ZARASHAW: If I may, it's a privacy  
5 preserving ORM layer.

6 MR. PATEL: Yeah. That's a good way to  
7 frame it.

8 MR. ZARASHAW: It's an ORM layer on top of  
9 a very raw data storage that is the graph database  
10 provided by TAO. TAO itself would not implement the  
11 privacy checks and the conflict logic of, let's say,  
12 if you make a comment on a post, understanding that  
13 you, as the author, are able to delete that comment  
14 and the author of the original post has the right to  
15 delete the comment, but no one else has the right to  
16 delete that comment. There are a tremendous number  
17 of rules like this all over every Facebook product  
18 and surface that it would be very error prone to  
19 have developers implement repeatedly.

20 So it's an ORM layer that lets you query  
21 the data in TAO and write data to TAO, but  
22 implementing all the privacy rules underneath. So  
23 no individual developer is not working on it.

24 MR. GARRIE: Got it. That makes sense.  
25 All right. That makes sense.

1 Okay. So then is Node the same?

2 MR. PATEL: Yeah. Correct. Node is the  
3 same. It's just a different language.

4 MR. GARRIE: It's for all the other ones, I  
5 guess. All right.

6 So then does Plaintiff have any questions?

7 MS. WEAVER: I think we're okay.

8 MR. GARRIE: So then MySQL, what  
9 associations -- well, what associations or objects  
10 from TAO does MySQL store? I didn't quite  
11 understand that.

12 MR. ZARASHAW: All of them. It is under --  
13 it is the backing store for TAO and we would need to  
14 get a TAO expert into the room for the details, but  
15 for every -- every piece of data that is stored in  
16 TAO is actually being stored in one of thousands and  
17 thousands of MySQL servers distributed globally and  
18 usually on more than one of them, and any schema  
19 updates to TAO match to some schema changes on those  
20 MySQL servers. So those MySQL servers are not  
21 generally used directly and they are used through  
22 TAO.

23 MR. GARRIE: I really like that ORM layer  
24 explanation. That cleared up a lot.

25 All right. So MySQL.

1 All right. So then [REDACTED] is there a  
2 type of data stored in [REDACTED] Is it binary like  
3 images or text like JSON objects?

4 MR. ZARASHAW: It's meant for -- so TAO is  
5 not efficient at storing large chunks of anything,  
6 binary or text. So once you get over a certain  
7 limit that I don't remember what the latest limit of  
8 MySQL is, last I saw it was 64K, MySQL performs very  
9 poorly.

10 The point of [REDACTED]  
[REDACTED]

12 MR. GARRIE: So then if there's JSON data  
13 in [REDACTED] can it contain user data?

14 MR. ZARASHAW: [REDACTED] [REDACTED]  
[REDACTED]  
[REDACTED]  
16 [REDACTED] [REDACTED]  
[REDACTED]

18 MR. GARRIE: Right. And then it would be  
19 [REDACTED] Got it.

20 MR. ZARASHAW: [REDACTED]  
[REDACTED] and  
22 the developer would not have to think about it.

23 MR. GARRIE: So is there a list of [REDACTED]  
[REDACTED]  
[REDACTED]

1 MR. ZARASHAW: I'm out of my depth here.  
2 Mayur?

3 MR. PATEL: We would have to go look at the  
4 configuration and see if that's possible to pull up.  
5 Again, that's out of my expertise as well.

6 MR. GARRIE: Okay. Zippy. The Zippy  
7 database and Akkio, do you know what associations or  
8 objects from TAO that Zippy or Akkio stores?

9 MR. ZARASHAW: I do not know.

10 MR. PATEL: I do not know. We generally  
11 understand how Zippy operates.

12 MR. GARRIE: Is Akkio used for -- it says  
13 it flips the datasets into units with strong  
14 locality, meaning geographical locality?

15 MR. ZARASHAW: Data center locality in this  
16 case.

17 One of the challenges where we say is we  
18 have many data centers globally, and any time a  
19 transaction needs to be replicated to all data  
20 centers, it's far more extensive than it would be  
21 locally.

22 MR. GARRIE: So then that's geographically  
23 placed close to where there, most likely, would be  
24 access. Okay. Makes sense.

25 MR. ZARASHAW: Exactly. What I mean by



1 extensively is latency. It takes a long time to  
2 commit or write across the continent; but if you  
3 only care about coming back to the same data center,  
4 it would create a better user experience.

5 MR. GARRIE: And I read through the  
6 documentation. Memcache made sense to me. I got  
7 that.

8 Laser, do Laser indexes point at the Hive  
9 and make the Hive query?

10 MR. ZARASHAW: It's the other way around.  
11 So a job will run. So Dataswarm jobs can have a  
12 source and a target. So a source for Dataswarm in  
13 this case would be Hive. It would run a set of  
14 queries or custom code and the target would be  
15 Laser. So that job would then populate a Laser tier  
16 with the results of a Hive query that could be --  
17 the result being potentially quite large.

18 There are other ways to populate Laser as  
19 well. It does not have to be from Hive. It could  
20 be done from a shelf script if the developer chooses  
21 to do so. It could be done from any code that's  
22 written, but the common use case would be run an  
23 extensive query in Hive and put the results in Laser  
24 for very fast lookups of specific rows.

25 MR. GARRIE: And Laser has storage

1 basically.

2 MR. ZARASHAW: It's a cache. It's a cache  
3 for storing a precomputation of a large dataset when  
4 you only ever want to access a small subset of that  
5 dataset.

6 Hive doesn't really offer the ability to  
7 access one row. Every operation in Hive ends up  
8 touching the entire table and this is one of the  
9 challenges we run into any time we are looking for  
10 user data in Hive. There's no way in Hive to just  
11 quickly grab a subset of the table in Hive. Every  
12 single lookup in Hive is a scan of the entire file,  
13 of the entire table. It involves the same amount of  
14 time to look at all of the data in Laser as to just  
15 one row.

16 MS. RING: Special Master Garrie, if I may,  
17 actually I forgot that on a break it was asked, they  
18 wanted to clarify something that had been talked  
19 about about user IDs and being in tables in Hive and  
20 what it would take to find the data in the tables in  
21 Hive that have the user ID.

22 I think, Mike and Mayur, you guys wanted to  
23 clarify.

24 MR. CLARK: Well, I think it was in  
25 connection to Laser. Like Laser has very specific

1 mapping and, as Eugene just walked through, that's  
2 why that one-to-one mapping makes it much more  
3 efficient.

4           There is the challenge of when we go to  
5 look for things broadly in Hive across tables and  
6 where we don't have that preexisting application or  
7 dataset mapping, that finding it, and I used the  
8 analogy earlier, but not just looking in one phone  
9 book or a couple phone books, it's the many  
10 warehouses of phone books with millions and millions  
11 of them trying to find the same info.

12           MR. GARRIE: So I get that.

13           You guys have five million -- you run  
14 through five million every night, every day. On  
15 one day you ran through five million. So it's not  
16 like you don't have the computational capacity to  
17 run it.

18           MR. CLARK: But that's only the  
19 computational capacity to map it.

20           What is the challenge is most of it's not  
21 actually in warm storage, and so to actually --

22           MR. GARRIE: Well, there's a whole -- so  
23 there's a whole separate Hive conversation about how  
24 it goes --

25           MR. CLARK: Okay. Yeah.

1 MR. GARRIE: -- outside. There's a Hive  
2 specialist because there has been a prior  
3 conversation about different tables in Hive and how  
4 it's in cold storage and it's compressed and zipped  
5 and it's a pain to access and all of that. I  
6 assume nobody here has that domain specific  
7 expertise.

8 MR. PATEL: We don't have the expertise.  
9 We have the general understanding of --

10 MR. CLARK: Yes.

11 MR. PATEL: -- like cold storage and how it  
12 operates.

13 MR. GARRIE: I've read -- I mean, there's a  
14 lot of previous advocacy around it.

15 MR. CLARK: And the only thing I saw  
16 missing in the advocacy was just the amount of time  
17 to then manually map all that.

18 MR. GARRIE: Oh, no. They explained that.

19 MR. CLARK: Okay. I didn't know. The  
20 stuff I read didn't have that. So apologies,  
21 Special Master Garrie.

22 MR. GARRIE: They probably spared you.  
23 There was a whole other path that was resolved  
24 around tables in Hive to tables in Hive and so on.  
25 So I'm fairly familiar with the amount of cold

1 storage sharing on a daily basis and how that's  
2 zipped and you have to unzip it and query it.

3 MS. RING: Special Master Garrie, look, I'm  
4 sorry to interject again, but we want to really help  
5 you understand this and I think you're thinking of a  
6 different issue. I've read all that documentation  
7 also.

8 When they were telling me this on a break,  
9 this is a distinct issue, which after it's pulled  
10 then it has to be reviewed, which I don't think  
11 we've addressed in prior submissions.

12 But if you don't want to hear about it, I  
13 mean, it is an important part of the process for  
14 being able to actually -- I mean, you know, Mike,  
15 this is --

16 MR. CLARK: It's the longest amount of  
17 time. Like it's -- it could take X amount of time  
18 to pull it, then it's X time to review it because it  
19 is manual because there isn't just a simple way to  
20 go technically audit what each of those tables were  
21 to then know how they were mapped.

22 MR. GARRIE: That's one big for loop and  
23 then someone has to put eyes on the output.

24 MR. PATEL: It's not just one big for loop.  
25 Right.

1 MR. GARRIE: I'm generalizing, but I get  
2 it. I get it.

3 MR. PATEL: It's a for loop with a whole  
4 bunch of technical constraints around, you know, how  
5 much cold storage we can move to warm on a single  
6 day. It's not that we can move our entire cold  
7 storage into warm storage.

8 MR. GARRIE: No. I get that. It's only  
9 one day.

10 MR. PATEL: Yes.

11 MR. GARRIE: So I get there's a lot of  
12 infrastructure issues because you move it to cold  
13 storage, and then you'd have to move it out of cold  
14 storage and then compress it, restore it, and that  
15 takes quite a bit of time before we even get to the  
16 search and all the other pieces of it.

17 MR. ZARASHAW: Sorry. Just one last thing  
18 on this point.

19 It's not just a matter of time. It's that  
20 we simply don't have enough hard drives to store  
21 additional -- to pull data out of cold storage  
22 without first deleting data that's already in warm  
23 storage that's being used for production.

24 MR. GARRIE: Yeah. I get it. I mean, that  
25 was delivered to me in the last train of thought we

1 had, which is why I granted the relief sought by  
2 Facebook at the time. So, yeah, I get that. I  
3 mean, 2.5 terabytes of data isn't chump change. So,  
4 you know, that's more storage than all of the data  
5 that's active litigation in the United States stored  
6 in relativity. Okay. So I fully grasp the  
7 magnitude of what we're talking about.

8 I mean, before we get down there, I will  
9 return to that conversation, Counsel Ring, and  
10 please bring it back up again. I just want to  
11 finish my train of thought to what we have.

12 I have specific questions about the product  
13 management process and specifically how that works.

14 But before I get to that, what -- so I  
15 think what -- you were saying what data does Laser  
16 index and I just want you to repeat it one more  
17 time.

18 MR. ZARASHAW: I think it might help if I  
19 contrast Memcache and Laser because they serve a  
20 very similar purpose, but used in two different  
21 ways.

22 I'd classify Memcache as a reactive cache  
23 in that you do an extensive computation and then you  
24 store the single result of that computation in  
25 Memcache so that if you within a time period have to

1 do that computation again, it is available for fast  
2 access.

3 Laser is more for cases where you cannot,  
4 on demand, do that computation. Instead, you do  
5 that computation, let's say, once a day in bulk on  
6 the entire Facebook user base, so the entire set of  
7 an object type and so on, and then you commingle  
8 all of it because the amount of time to do the  
9 computation for one item is almost the same as to do  
10 the computation for the entire universe.

11 So Laser is more of a proactive cache where  
12 we will do the extensive computations once in a day,  
13 we'll load it into a Memcache-like tier and allow  
14 randomized access to the sub-results of those  
15 computations.

16 MR. GARRIE: That makes sense. Because it  
17 says Laser is an indexing service, but it also has  
18 storage.

19 MR. ZARASHAW: Laser is effectively a set  
20 of Memcache-like machines plus infrastructure to  
21 run -- most of the time run a Hype query and load up  
22 the results. So the storage is ephemeral.

23 MR. GARRIE: Okay. So then how is Laser's  
24 use restricted to putting first-party only products  
25 if Graph API uses Laser to access the data?



1 MR. ZARASHAW: There is no direct access to  
2 Laser. So it would have to be through Graph API. I  
3 don't know which specific APIs use Laser, but it  
4 would be similar to Memcache where Memcache would be  
5 used when a specific API can't be fulfilled quickly  
6 by getting MySQL or TAO directly.

7 MR. GARRIE: And then are there private  
8 internal-only Graph APIs between the two?

9 MR. ZARASHAW: It would not be a Graph API.  
10 Graph APIs are externally available. So Laser  
11 offers up a drift service for any --

12 MR. GARRIE: Right. If Laser is restricted  
13 to supporting first-party only products and if  
14 Graph API uses Laser to access the data and it's not  
15 done, how does that work?

16 MR. ZARASHAW: So -- I'm sorry.

17 MR. CLARK: Sorry.

18 Mayur, didn't you have a good example of  
19 like the OFA process, for instance? Maybe going  
20 through a concrete example might actually be helpful  
21 in this case.

22 MR. PATEL: I am actually confused with the  
23 question. So if the question can be repeated.

24 MR. GARRIE: Yeah. So let me explain as I  
25 understood it and then maybe I'm completely wrong,

1 which is possible.

2 So then Laser's use is restricted I thought  
3 to supporting first-party only products.

4 MR. ZARASHAW: That is not correct. We  
5 don't think about pieces of infrastructure as  
6 designed for first or third-party products. We have  
7 an API model for all third-party interactions  
8 happening through a set of external APIs, but then  
9 each API has product specific implementations  
10 because the API is very broad and many of the  
11 products we offer have their own APIs and then each  
12 API implementation may use different types of  
13 storage to accomplish the product outcomes of that  
14 API.

15 MR. GARRIE: Okay. So it doesn't really,  
16 frankly, matter. Okay. All right.

17 Question six, my question is more about the  
18 product process.

19 In the product process, how do these  
20 agreements -- how do agreements get made? I still  
21 haven't been able to figure this out.

22 Like, you have products and then they sign  
23 agreements with the customers for the products. So  
24 Netflix, Microsoft, whoever. You build a product.  
25 Like Facebook builds products and then there are

1 agreements with customers to use features of the  
2 products.

3 So what I'm trying to understand is what's  
4 the sequence of how that works and then if they're  
5 customer specific driven products or projects.

6 I just started reading over some of the  
7 agreements. So I'm not exactly sure.

8 MS. RING: So we really don't have someone  
9 on here who focuses on that. I mean, I think  
10 there's like a business tool and they're subject to  
11 standard terms.

12 MR. GARRIE: Well, no. That's my question.  
13 So that's my question.

14 You have products that you build and  
15 there's a set of -- how does this -- like I'm a  
16 business. I want to use your products.

17 MS. RING: So, I mean, look, we don't have  
18 anyone on this call that deals with that exactly.

19 I think if it's something -- if it's a  
20 business tool and it uses a standard app, you know,  
21 they just sign up and use it and it's subject to the  
22 standard terms.

23 I don't know for this specific. You know,  
24 we provided contracts for Netflix, Microsoft,  
25 YouTube. I don't know the process that's followed

1 and I can say who cares what I know, but we didn't  
2 bring anybody who knows that process.

3 MR. GARRIE: And all I'm trying to figure  
4 out is the pro- -- I get the products.

5 What I'm trying to figure out is these  
6 agreements, are the products -- I'm trying to figure  
7 out if there's a process for where products may be  
8 tailored, customized, whatever, for a particular  
9 customer and if there's a process, a product  
10 development process for that. You usually get a  
11 project program. Like usually there's a -- in most  
12 places there's a project manager. There's a product  
13 and there's a project manager assigned to build that  
14 out to meet the requirements to tailor that specific  
15 need or whatever.

16 Just at a high level for those products, is  
17 there a standardized process that's followed?

18 MR. CLARK: At just the highest level, like  
19 just to walk through this, there are the primary  
20 products. There is the third-party developer  
21 platform, the business products that we've already  
22 spoken about. For any changes to any of those, they  
23 go through both the API XFN that I referred to in  
24 addition to privacy review for those checkpoints.

25 MR. GARRIE: But there are agreements that

1 are made with clients like Netflix or with  
2 Microsoft, for example.

3 MR. CLARK: And that's where I can talk to  
4 the example overall. And when there are specific  
5 individual products, I am not an expert in that.

6 MR. GARRIE: But my question is, are they  
7 handled -- my -- it's just a broad question and  
8 that's all I have for question six is, is there a  
9 broad process for when a customer, a third party,  
10 Microsoft, Amazon, comes and says we want to do --  
11 we're a big boy sitting at the big boy table,  
12 right -- I have a six-year-old, so it's a  
13 commonly-used phrase in my house -- and they want  
14 certain things. Right. You know, I want my red  
15 cup. I want my whatever it is and there's a  
16 process, you know, that is followed.

17 And my question is at Facebook, what I'm  
18 trying to figure out, is there some process that's  
19 documented as to how this is done for -- you know,  
20 because it's involving -- like if you look at  
21 Microsoft, there's user information going back and  
22 forth all over the place.

23 Is there a process that's followed for the  
24 product that's documented to build it out generally?

25 MR. CLARK: They follow the same process.

1 They follow the same API XFN and privacy review.

2 There is a specialized business development  
3 team that works on the contracts side and with the  
4 product teams, but that is as far as this goes.

5 They follow the same types of constraints and  
6 controls and usage of data that all of the other  
7 teams do as part of the overall program.

8 MR. MITCHELL: There are no exceptions to  
9 our standard development process or deviations.

10 MR. GARRIE: That's basically what I'm  
11 getting at.

12 MR. MITCHELL: Yeah. We have a common  
13 development process that is utilized by, you know,  
14 sort of all products and they all go through the  
15 same privacy review and, you know, they'll go  
16 through all the gates we've discussed at various  
17 points over the course of today and there's no  
18 exception because you're Microsoft or Netflix and  
19 you're a big boy.

20 MR. GARRIE: That's why I'm asking because  
21 there's a whole discussion of private APIs and  
22 everything else, but it follows -- today it follows  
23 that process.

24 When did this process get in place?

25 MR. CLARK: I can't speak to the exact

1 date.

2 MR. MITCHELL: It's before my time.

3 MR. GARRIE: Well, we can work backwards  
4 from 2020, but it would be helpful to know when that  
5 process came into place given the class period is  
6 rather -- because that applies to new things, not  
7 old things; right? So one of the questions is --  
8 but that's not for now, frankly.

9 So let me keep going in the interest of  
10 everybody that's here and the engineering resources  
11 that are here.

12 Question seven, how is Hive used? Is Hive  
13 used as a data warehouse while building product  
14 experiences? I couldn't figure it out from the  
15 description.

16 MR. ZARASHAW: What do you mean by while  
17 building product experiences? Like for what part of  
18 the product experience do you mean?

19 MR. GARRIE: I mean start to finish.

20 MR. ZARASHAW: Most commonly Hive would be  
21 used for the analytics part, like the data warehouse  
22 for how a project is doing rather than for  
23 fulfilling the primary project itself, more to  
24 inform us about whether the product is actually  
25 working, if it's useful, and so on.

1 MR. GARRIE: I mean, that makes sense given  
2 the amount of data that's being collected.

3 I guess now that I understand a little  
4 better, is a query to build experiences assuming  
5 you're using, like, Laser? Is that how that works?

6 MR. ZARASHAW: I'm not sure I understand  
7 the question.

8 MR. GARRIE: So how are you querying -- I  
9 think I --

10 Never mind. You've already answered the  
11 question. Because Laser is like Memcache. It's  
12 indexed and it's targeted so you can then -- okay.

13 So then is there any clear delineation of  
14 any of the Hive data resources that are used in  
15 products? Is there a table key or something like  
16 that? Not for 100 percent.

17 MR. ZARASHAW: Not that I know. Mike?

18 MR. CLARK: No. And nothing comes directly  
19 to Hive. It would -- because it's just wait and  
20 see. It's like it's optimized for writes and not  
21 for reads and, from a product experience  
22 perspective, it would need to go through one of the  
23 other layers or some other transformation in order  
24 to be usable in a product experience.

25 MR. ZARASHAW: Actually that's a really



1 good point to call out, Mike.

2 Special Master Garrie, on a given -- as of  
3 about 8:00 a.m. or 9:00 a.m. for most tables, Hive  
4 contains the data for yesterday. It is not anywhere  
5 near real time. So at any point it takes somewhere  
6 between 8 and 12 hours depending on which table and  
7 sometimes 24 just to process the previous day's data  
8 and make it available for both queries. So it would  
9 not work for any production data use that would  
10 require us to read data back after writing it.

11 MR. CLARK: You, as a product engineer,  
12 would start with TAO and most of the time end with  
13 TAO.

14 MR. GARRIE: All right. So then -- okay.  
15 So then on Exhibit B, which, by the way, was  
16 helpful, the last row on Page 3 says other datasets,  
17 that Facebook occasionally acquires datasets to help  
18 inform different user models, blah, blah, blah.

19 Do you see where I'm at?

20 MR. PATEL: Sorry. Which page is this?

21 MR. GARRIE: Page 3, last row.

22 MS. RING: So, yeah. This is --

23 MR. MITCHELL: Page 41 in the pdf.

24 MR. GARRIE: Yeah.

25 MS. RING: Okay.

1 MR. CLARK: Found it.

2 MR. GARRIE: My question is, while census  
3 data is not user specific, implications derived from  
4 census data would be or is or can be.

5 So what are the data points that are  
6 stored? Like where are they stored?

7 MR. PATEL: Do you mean stored with the  
8 user at a user level granularity?

9 MR. GARRIE: Derivation from the census  
10 data. You're deriving data from the census data.  
11 Where is that data stored?

12 MR. PATEL: Which system are you talking  
13 about? Because then we'd probably need an admin  
14 engineer to describe like where that data  
15 eventually --

16 MR. ZARASHAW: It's like a lot of other  
17 situations where the life cycle is usually figuring  
18 out what use case we're trying to come up with.  
19 Then a product team decides a use case should exist  
20 in the world. They go build the use case, which  
21 may include acquiring a third-party dataset. Then  
22 for the use case they're trying to hit usually based  
23 on the read and write patterns, they'll pick the  
24 piece of infrastructure or system that makes sense  
25 for that use case. And so there's no general

1 principles for where it should be stored, other  
2 than what makes sense for the specific project.

3 MR. GARRIE: So I guess my follow-up  
4 question to that is, are these implications included  
5 in the DYI file?

6 MR. PATEL: The implica- -- if you're  
7 talking about the ads targeting criteria, they are  
8 in DYI.

9 MR. GARRIE: I'm talking about any of the  
10 data. Well, no. You said ads. I'm saying any of  
11 the --

12 What I'm being very clear about is derived  
13 from the census data because you say there's, you  
14 know, other datasets and you reference the census  
15 data. I'm saying you can derive data from that  
16 information and my question is, is the derived  
17 information that you get from that stored or made  
18 available in the DYI file?

19 MS. RING: Excuse me, Special Master  
20 Garrie.

21 I mean, Mr. Patel, this is -- we talked.  
22 This is -- this column is -- this column is  
23 referring to something very specific, which is ad  
24 targeting options based on income. Those are in  
25 DYI.

1           So, Mr. Patel, I mean, can you just  
2     describe -- I think it would be helpful to describe  
3     this is just an ad targeting option that is created  
4     using census data we get and that is in the DYI file  
5     and people can X out of it, et cetera.

6           MR. PATEL: I mean, you have transparency  
7     and control in your ads preferences as well into  
8     this.

9           MS. RING: Maybe it's actually --

10          MR. GARRIE: Well, let me make it simpler.

11          Where might it be that I'm rich stored?  
12     Where is it stored? I mean, I guess that's what I'm  
13     trying to figure out. Right.

14          MS. RING: In your DYI file. I looked mine  
15     up and I'm apparently not rich. My ad targeting  
16     option was actually for a low -- like a low  
17     percentage based on my zip code, I guess. So --

18          MR. PATEL: Yeah. Special Master Garrie,  
19     I don't know that answer because it's the ads  
20     targeting product team that would know exactly where  
21     this is stored. Like, I can just speculate where I  
22     would store it, but I cannot tell you for sure.

23          MR. GARRIE: And then because one thing I  
24     noticed looking at the APIs, you require data to  
25     be -- just returning back to a prior conversation.

1           In some of the API functions, you require  
2     it to be in there and you don't necessarily produce  
3     that data in the DYI file every single time and  
4     there's -- there are specific elements, but do you  
5     have data elements? Like, for example, the  
6     Luluemon, you require there to be a value put in  
7     there, but it doesn't mean that it will necessarily  
8     be returned to the DYI file at least based on the  
9     documentation that I looked at.

10           MR. CLARK: I believe, and I'll go back and  
11     reiterate what we'd stated before. And you broke up  
12     for just a moment in the question, so I want to make  
13     sure I'm answering this right.

14           The only required fields are the, you know,  
15     time and date stamp that comes in, the event type,  
16     and we share that back. There are custom fields  
17     that the third party, the business product user  
18     company, individual can apply on their own and that  
19     is the content that to us is not directly mapped.  
20     If they map a page to you, if they map an initiate  
21     checkout, those are things that are specifically  
22     mapped and meaningful.

23           MR. GARRIE: Okay. I get it.

24           So then what are the other sources of data  
25     that Facebook draws implications from?

1 MS. RING: I think you cut out for me, at  
2 least. I couldn't hear. It sounded like there was  
3 an echo.

4 MR. GARRIE: What are other sources of data  
5 that Facebook draws implications from?

6 MR. ZARASHAW: I think we would need to get  
7 somebody from the ads targeting team to answer  
8 that.

9 MR. GARRIE: Yeah. It sounds like the ads  
10 targeting team will be -- they will be able to speak  
11 to the use of cookies that aren't the user, the  
12 identifiers, but are different cookie types as well  
13 as this area. But, otherwise.

14 And then my last question, on Exhibit C can  
15 you just explain to me the last two sentences in the  
16 opening paragraph?

17 MR. CLARK: On the spacer page that says  
18 Exhibit C?

19 MR. GARRIE: Yeah. In the opening  
20 paragraph on Exhibit C. I just didn't understand  
21 it.

22 MR. CLARK: Is that a question for us, the  
23 engineers, or is that a question for the legal  
24 team?

25 MR. GARRIE: The legal. I mean, whoever is

1 best to answer it. I have a feeling the legal team  
2 may, but I don't know.

3 MS. RING: I'm sorry. Can you repeat the  
4 question? I didn't -- I couldn't -- I'm going out.  
5 I don't know if it's mine.

6 MR. GARRIE: The last two sentences in the  
7 opening paragraph.

8 MS. RING: Of what page?

9 MR. GARRIE: Exhibit C.

10 MR. MITCHELL: Page 42.

11 MR. CLARK: Is it the thing that says, "On  
12 3 DYI includes data that as a factual matter"? Is  
13 that what you're referring to?

14 MR. GARRIE: Yeah.

15 MS. RING: What does that mean?

16 MR. GARRIE: Yeah.

17 MS. RING: It just means that inference is  
18 not a defined term, and so it has legal meanings in  
19 certain contexts and we're just talking about as a  
20 factual matter.

21 Like in the table we just looked at at the  
22 bottom, those are ad targeting options. Those are  
23 in DYI. They're derived from census data and zip  
24 code and that's it. And so it's just to say like  
25 that's just derived data.

1 MR. GARRIE: Okay. And then we need to  
2 talk to someone on the ads team, it sounds like, to  
3 determine what else besides census and how that  
4 flows.

5 MS. RING: For ad targeting? Like the ad  
6 targeting?

7 MR. GARRIE: Yeah. Besides census in  
8 Exhibit B.

9 MS. RING: Yeah. I mean, just to be -- you  
10 mean -- you mean ad targeting, all those options are  
11 in DYI, but you want to know where they're stored  
12 within --

13 MR. GARRIE: But, more importantly -- yeah.  
14 But, more importantly, what is stored when they're  
15 providing -- what data do they get and then what  
16 data is going in the DYI? Is it one-to-one or is it  
17 a --

18 MS. RING: It's based -- the ad targeting  
19 options are derived from on-platform activity, and  
20 then the targeting option goes into DYI. I'm sure  
21 it's stored somewhere else for use in advertising,  
22 but is that the -- so that's the question?

23 I think and, Mayur, you just don't know the  
24 answer to that because you're not on the ads team.

25 MR. PATEL: Yeah. I don't know where it's



1 stored, which specific system.

2 MR. GARRIE: And I have a couple other  
3 questions there, but I need to think over the  
4 answers that were given before, but I can give them  
5 in writing. I have a couple scenarios; but rather  
6 than go through the scenarios now, what I'm  
7 thinking I'll do is I'll provide Facebook a set of  
8 scenarios to explain the specific outputs of where  
9 data goes.

10 MS. RING: I'm sorry. You cut out at the  
11 end. What did you say? What was the end part of  
12 what you said?

13 MR. GARRIE: I'm going to provide a set of  
14 scenarios, in all likelihood. I had them written  
15 here. Like Plaintiff suggested some, but I actually  
16 have my own. But I don't think it -- it would  
17 probably take all day to go through the scenarios.

18 So rather than taking everybody's time to  
19 go through them, I was thinking about putting the  
20 scenarios in writing and letting you guys provide  
21 written answers explaining, you know, what data is  
22 being collected and where is it going rather than  
23 doing it now.

24 MS. RING: That would make it interesting.

25 MR. GARRIE: It wasn't nearly as -- yeah.

1 But, anyways, that's what I'm thinking we'll do  
2 rather than doing it in real time right now trying  
3 to figure it out, but give you guys a certain time  
4 to just review and then give the written answers.

5 Has anybody ever run a disk between TAO and  
6 the DYI file?

7 MR. PATEL: I have not. I don't -- I don't  
8 think we can. Like, they're two different schemas.

9 MR. GARRIE: You'd have to map them, but  
10 you could.

11 All right. I think that covers most of my  
12 questions for today, and I think my preference is  
13 to provide the written and additional follow-on  
14 questions rather than continuing forward.

15 I may have one last -- I need to review  
16 what was covered today. We covered a lot, but  
17 there's certainly some further conversations about  
18 the Hive and the advertising, the cookies and how  
19 those -- like you have a range of cookies, how are  
20 they incorporated and used into the system.

21 I understand the systems part now. Now I  
22 need the data piece on that. Nobody here can speak  
23 to DATR cookies or any of the other cookies that are  
24 used?

25 MR. PATEL: Sorry. Which cookies are you

1 referring to specifically?

2 MR. GARRIE: There's on-platform Facebook  
3 cookies, and there's third-party cookies, and then  
4 there are subsets of those as well. I mean, we can  
5 run through them if you want, but I didn't think  
6 that -- the DATR cookie.

7 MR. PATEL: Yeah. The DATR cookie is  
8 primarily used for security purposes. It's to  
9 understand whether or not you have previously logged  
10 in and the primary purpose to whether or not you've  
11 primarily logged in on that browser. So it  
12 identifies a browser. And so if you try to log in  
13 again -- sorry. If you try to log in the first time  
14 with a new DATR cookie, we may checkpoint you there  
15 and ask for a confirmation through email or phone in  
16 order to continue to log you in. And then if you  
17 log out and log back in with the same DATR cookie,  
18 we may not ask you again. So that's the reason for  
19 the first-party cookie that we have.

20 MR. GARRIE: No. My question was  
21 different.

22 Why is it in the DYI file and not other  
23 cookies? I understood what it did. We spoke  
24 previously about it.

25 MR. PATEL: Oh.

1 MR. GARRIE: My question was more, why is  
2 that cookie put in the DYI file and who's making  
3 the decision about what cookies are going into the  
4 DYI file, what are all the cookies you're capturing  
5 and storing in the Hive, how are they associated  
6 with the users, and how are they being shared and  
7 aggregated to third parties?

8 If you guys can speak to that, we can get  
9 into it; but if not, I won't recommend it.

10 MR. PATEL: No.

11 MR. CLARK: No. That is one we would need  
12 to get some additional folks engaged with.

13 MR. GARRIE: No. I mean, so I got what  
14 it's used for. I'm just trying to figure out why  
15 some cookies are. Like, there's a lot of cookies  
16 that exist in the world, like a lot, and a user will  
17 have many. And so how are they -- what and how is  
18 being captured and utilized and passed and shared  
19 back and forth and utilized?

20 So I figured I'd table that. So that I  
21 have follow-on. I may have follow-on at one single  
22 hearing to cover the Hive and those issues, and then  
23 send you the scenarios is what I'm thinking.

24 But before I wrap it up, Plaintiff, do you  
25 have any questions?

1 MS. WEAVER: No. I think this has been  
2 very, very helpful and I think it would be useful  
3 for us to read the transcript and be thoughtful  
4 before following up.

5 But I wanted to thank everybody, thank  
6 Facebook and all of the engineers who came to try to  
7 educate us today.

8 MR. GARRIE: Yeah. So thank you to  
9 Facebook. Thank you to all four engineers for  
10 taking time out of your very busy day. I know the  
11 lawyers have their jobs, but your job is not to,  
12 per se, participate in these proceedings on the  
13 regular. So I really appreciate you taking the  
14 time out of your day and the time to prepare to  
15 answer these questions and I wanted to thank you for  
16 making yourselves available. It was extremely  
17 informative and helpful. And I hope we don't talk  
18 again; but if we do, then hopefully it will be as  
19 instructive and as useful as this time. So thank  
20 you all very much.

21 And do we need to do anything,  
22 Court Reporter, before we dismiss everybody off the  
23 record or anything like that?

24 MS. RING: I think Mr. Solanki is trying to  
25 say something.

1 MR. SOLANKI: Sorry. I may have missed  
2 this.

3 We had talked, I think, in the last hearing  
4 about just designating it as highly confidential  
5 given the amount of information that's been  
6 disclosed about our systems. So you may have done  
7 that already. But if not --

8 MR. GARRIE: I did not. I did not.

9 So he's 100 percent correct. This is  
10 highly confidential and should be designated as  
11 such and treated accordingly.

12 THE COURT REPORTER: Okay.

13 MS. RING: May I add one last thing,  
14 Special Master Garrie, while we have everyone here,  
15 including Plaintiff?

16 On the turnaround time for these scenarios  
17 or additional questions, we'd appreciate as much  
18 time as we can get. It's hard to turn them around  
19 and give good and full answers when we don't have  
20 very much time.

21 MR. GARRIE: I fully acknowledge and hear  
22 you. When I next talk with Judge Chhabria, I'll  
23 communicate accordingly. But I agree with you and  
24 appreciate and understand how much work and effort  
25 is involved here and I will attempt to convey the

1 same, as appropriate.

2 MS. RING: Thank you.

3 Look, the desire is to have a helpful  
4 response. It's easier to do that with a little more  
5 time. That's all. So thank you.

6 MR. GARRIE: Yeah. Noted for the record,  
7 certainly.

8 And we'll go off the record here, unless  
9 anybody has anything else to say.

10 \*\*\*

11 (Whereupon, the hearing ended at 2:28 p.m. PST)

12

13

14

15

16

17

18

19

20

21

22

23

24

25


CERTIFICATION OF COURT REPORTER

FEDERAL JURAT

I, the undersigned, a Certified Shorthand Reporter of the State of California do hereby certify:

That the foregoing proceedings were taken before me at the time and place herein set forth; that any witnesses in the foregoing proceedings, prior to testifying, were placed under oath; that a verbatim record of the proceedings was made by me using machine shorthand which was thereafter transcribed under my direction; further, that the foregoing is an accurate transcription thereof. I further certify that I am neither financially interested in the action nor a relative or employee of any attorney of any of the parties.

IN WITNESS WHEREOF, I have this date subscribed my name: Date: March 14, 2022.



Michelle Milan Fulmer

CSR 6942, RPR, CRR, CRC



[&amp; - add]

<b>&amp;</b>	<b>2843</b> 1:5 2:5	<b>98101</b> 3:11	<b>account</b> 110:12
<b>&amp;</b> 3:4,16,20	<b>29</b> 14:16,17,18	<b>9:00</b> 145:3	114:6
<b>0</b>	<b>2:28</b> 5:2 159:11	<b>a</b>	<b>accuracy</b> 92:8
<b>02843</b> 1:6 2:6	<b>3</b>	<b>a.m.</b> 2:12 5:2 64:1	<b>accurate</b> 31:18
<b>1</b>	<b>3</b> 13:16 93:7	64:2 145:3,3	56:1 62:22 74:15
<b>1</b> 4:3 13:13 42:19	145:16,21 151:12	<b>ability</b> 41:8 130:6	89:22 160:14
49:1 61:1 64:5	<b>3000</b> 3:21	<b>able</b> 15:25 23:6	<b>achieve</b> 52:18
<b>1.0</b> 67:7	<b>3200</b> 3:11	51:19 53:2 87:14	<b>acknowledge</b>
<b>10</b> 101:3 117:7	<b>393-8247</b> 3:22	87:19,20,23 92:5	158:21
<b>100</b> 11:17 92:8	<b>4</b>	108:22 116:17	<b>acquires</b> 145:17
144:16 158:9	<b>41</b> 145:23	119:21 125:13	<b>acquiring</b> 146:21
<b>10:20</b> 2:12 5:2	<b>415</b> 3:6,22	133:14 138:21	<b>acronym</b> 65:1
<b>11:36</b> 64:1	<b>42</b> 151:10	150:10	74:19 78:7
<b>11:54</b> 64:2	<b>445-4003</b> 3:6	<b>absolute</b> 19:3,4	<b>action</b> 80:18
<b>12</b> 145:6	<b>45</b> 5:6	<b>accelerated</b> 6:23	160:16
<b>1201</b> 3:11	<b>5</b>	<b>acceptable</b> 10:23	<b>active</b> 49:18 135:5
<b>12th</b> 3:5	<b>513</b> 24:8	82:17	<b>activities</b> 13:16
<b>14</b> 160:19	<b>555</b> 3:5,21	<b>accepted</b> 83:20	64:12,15
<b>149</b> 13:3 18:7	<b>6</b>	<b>access</b> 14:15,19,22	<b>activity</b> 13:14,17
<b>15030</b> 160:23	<b>623-1900</b> 3:12	19:7 37:14 38:25	16:7,9,10,12,16,16
<b>16</b> 25:13,14	<b>64</b> 105:14 106:1	41:18 42:6,8	16:20,21,25 17:2,6
<b>1600</b> 3:5	107:14,25 109:6	51:18 53:4 68:6	18:2,10,21,22
<b>18</b> 1:6 2:6	<b>64k</b> 127:8	84:10 90:12	19:15 20:9,13
<b>1881</b> 3:17	<b>650</b> 3:18	106:20,21,23	21:6 22:11 27:5
<b>1:09</b> 117:9	<b>6942</b> 1:18 160:25	107:4 108:8,9,12	39:9 49:5,12
<b>1:31</b> 117:10	<b>8</b>	125:1,3 128:24	50:10 51:20,21
<b>2</b>	<b>8</b> 145:6	130:4,7 132:5	61:5 68:25 74:20
<b>2</b> 13:14 38:17	<b>849-5348</b> 3:18	136:2,14,25 137:1	76:3 78:8 85:16
60:25 64:6 86:20	<b>8:00</b> 145:3	137:14	86:5 88:14,15,21
<b>2.0</b> 67:7	<b>9</b>	<b>accessed</b> 92:22	89:13,24 90:2,2
<b>2.5</b> 135:3	<b>9</b> 1:15 2:13 5:1	108:16 121:5,21	103:16 114:21
<b>20</b> 85:22	<b>90</b> 112:15	122:8	120:21 152:19
<b>2007</b> 66:24	<b>94025</b> 4:4	<b>accessible</b> 24:1	<b>actual</b> 39:19
<b>2020</b> 143:4	<b>94105</b> 3:22	127:16	<b>ad</b> 70:13,13,17
<b>2022</b> 1:15 2:13 5:1	<b>94304-1211</b> 3:17	<b>accessing</b> 122:19	72:21 81:12,14,22
160:19	<b>94607</b> 3:6	124:21	81:24 82:6 147:23
<b>206</b> 3:12	<b>95</b> 19:3 65:18,21	<b>accidentally</b> 95:25	148:3,15 151:22
<b>20th</b> 56:9,10	65:25 110:5	<b>accommodating</b>	152:5,5,10,18
<b>24</b> 145:7		6:23	<b>adavis</b> 3:7
		<b>accomplish</b> 102:12	<b>add</b> 9:1 11:6,7
		138:13	12:22 20:13,17

[add - appropriate]

21:5 27:14 28:2 28:21 32:10 44:6 81:17 94:9 95:22 97:8,17,19,21 120:25 122:21 158:13 <b>added</b> 21:1,13 27:4 94:7 95:12 98:4 <b>adding</b> 20:8 <b>addition</b> 72:25 73:4 140:24 <b>additional</b> 12:17 28:11 61:14 72:17 134:21 154:13 156:12 158:17 <b>additions</b> 73:1 <b>address</b> 10:14 49:14 50:15 53:7 88:22 89:18 <b>addressed</b> 133:11 <b>addresses</b> 49:14 <b>adjourn</b> 63:9 <b>admin</b> 70:5 146:13 <b>admitting</b> 10:8 <b>ads</b> 64:19 65:1 81:16,23 88:13,18 88:19 147:7,10 148:7,19 150:7,9 152:2,24 <b>advertiser</b> 70:7 79:17 80:17 81:24 83:12 85:4 86:2 97:14 <b>advertiser's</b> 86:7 <b>advertisers</b> 49:13 78:25 82:24 83:6 85:16 95:20 <b>advertises</b> 70:7 <b>advertising</b> 152:21 154:18	<b>advocacy</b> 132:14 132:16 <b>aggregate</b> 75:6 <b>aggregated</b> 17:13 51:25 156:7 <b>aggregates</b> 84:8 <b>aggregator</b> 52:10 <b>ago</b> 66:4,6 <b>agree</b> 13:2,6 26:18 29:5 31:6 42:10 66:12,13 72:9 158:23 <b>agreement</b> 66:11 <b>agreements</b> 67:19 138:20,20,23 139:1,7 140:6,25 <b>ahead</b> 20:6 37:3,4 37:5 67:2,3 <b>akkio</b> 128:7,8,12 <b>algorithmic</b> 89:10 <b>allow</b> 136:13 <b>allowed</b> 38:25 39:3,12 124:10,11 124:12 <b>allowing</b> 15:1 <b>allows</b> 15:24 <b>alphanumeric</b> 118:12,13 120:3,4 <b>alto</b> 3:17 <b>amazing</b> 28:23 <b>amazon</b> 141:10 <b>ambiguity</b> 13:10 13:22 62:14 <b>america</b> 24:9 <b>amount</b> 55:23 67:23 102:9 130:13 132:16,25 133:16,17 136:8 144:2 158:5 <b>analogy</b> 65:24 131:8	<b>analysis</b> 39:10 57:9 89:11 90:3 91:3 93:2 106:18 <b>analytics</b> 143:21 <b>android</b> 24:8 <b>anne</b> 3:5 <b>annotating</b> 114:24 <b>annoying</b> 91:17 <b>ansi</b> 34:9,17 <b>answer</b> 7:3,4 11:3 11:14 17:22,23,25 18:17 32:16,18 33:9 44:24 47:4 48:2 59:5 61:11 61:13,14 73:8 74:15 78:21 83:13 84:22 87:3 94:4 100:13,19 101:10 102:21 105:3 111:11 117:2 118:19 148:19 150:7 151:1 152:24 157:15 <b>answered</b> 32:1,12 144:10 <b>answering</b> 61:18 71:12 101:8 149:13 <b>answers</b> 7:6 11:4 14:1 46:12 62:12 75:11 82:23 104:4 153:4,21 154:4 158:19 <b>anybody</b> 5:10 20:16 39:7 55:8 61:18 68:23 96:21 97:17 115:10 118:20 127:15 140:2 154:5 159:9 <b>anyway</b> 29:14	<b>anyways</b> 127:19 154:1 <b>api</b> 24:5 49:25 64:19,25 65:1,5 67:13 69:7 72:20 72:25 73:13,21,22 73:23 78:25 80:8 97:13 136:25 137:2,5,9,14 138:7 138:9,10,12,14 140:23 142:1 149:1 <b>apis</b> 67:5,8 68:24 72:10,17,22 73:10 99:9 137:3,8,10 138:8,11 142:21 148:24 <b>apologies</b> 65:1 132:20 <b>apologize</b> 5:12 8:3 44:7 <b>app</b> 64:22,24 65:5 68:10,12 86:5,7 97:15 139:20 <b>apparently</b> 148:15 <b>appearances</b> 3:1 <b>apples</b> 14:17,17 <b>application</b> 131:6 <b>applied</b> 108:5 <b>applies</b> 143:6 <b>apply</b> 123:12 149:18 <b>appreciate</b> 12:24 19:11 50:6 63:9 113:12 157:13 158:17,24 <b>approach</b> 12:14 113:13 <b>approaching</b> 13:3 <b>appropriate</b> 6:8 7:17,24 54:17
--	---	--	--

[appropriate - best]

159:1 <b>approval</b> 28:16 <b>approve</b> 41:14 <b>approved</b> 29:21 29:24 30:1 <b>apps</b> 67:15,16 <b>architected</b> 54:11 <b>archive</b> 16:21 25:18 <b>area</b> 150:13 <b>areas</b> 10:22 11:1 21:6 22:18 <b>arguably</b> 47:22 <b>asked</b> 7:5 48:11 58:5 85:11 95:15 105:3 130:17 <b>asking</b> 6:12 7:12 8:14 18:2 43:8 44:15,16 60:11 83:21 85:10 105:5 109:16 111:21 120:3 142:20 <b>asks</b> 61:12 <b>aspect</b> 91:13 <b>assign</b> 63:17 <b>assigned</b> 140:13 <b>assignment</b> 106:5 <b>associate</b> 58:24 <b>associated</b> 22:18 51:7 76:15 79:16 85:15 104:24 107:19 110:18,20 110:21 123:6 156:5 <b>association</b> 75:23 75:24 103:20 <b>associations</b> 117:14,25 119:5 126:9,9 127:24 128:7	<b>associative</b> 23:17 <b>assume</b> 18:13 22:6 25:16 32:25 36:23 41:5 45:3 61:23 67:19 74:2 76:23 86:23 90:18 99:25 100:2 111:11 124:3 132:6 <b>assuming</b> 39:6 41:4 42:16 62:15 100:8 144:4 <b>assumption</b> 31:5 62:18 74:15 <b>assumption's</b> 103:12 <b>attach</b> 23:16 <b>attached</b> 83:6 <b>attempt</b> 158:25 <b>attempting</b> 6:7 <b>attorney</b> 160:17 <b>attorneys</b> 1:10 <b>attribute</b> 70:12 <b>audience</b> 79:9 83:5,17 86:16 <b>audiences</b> 82:25 84:18,25 85:3,6,12 85:18,23 <b>audio</b> 5:14 <b>audit</b> 133:20 <b>auld</b> 3:4 <b>authentic</b> 26:7 <b>author</b> 124:7 125:13,14 <b>authorized</b> 39:8 <b>automated</b> 26:6 <b>automatically</b> 61:13 120:10 122:15 <b>available</b> 7:9 53:5 65:17 66:5,7,7 67:17,18 70:24	71:2,3,4,8,22 72:13,14,16,20 136:1 137:10 145:8 147:18 157:16 <b>avenue</b> 3:11 <b>avoid</b> 61:15 <b>aware</b> 10:15 19:16 44:23 73:12 <b>ayi</b> 15:2,8,9,10,12 15:13,15,24 16:15 16:22 17:1,1 18:23 20:8,13 21:6,20 27:4 31:10,14,24 56:5 56:14 80:1  <b>b</b>  <b>b</b> 145:15 152:8 <b>back</b> 13:19,20 26:9,12 29:15,17 36:21 37:11,13 38:7 54:4 62:1 64:1,3,4 68:25 69:4 70:12 74:22 76:15 77:12 79:2 81:13 83:3 89:4 89:12 93:9 94:10 94:22 102:20 105:15 108:9 110:20 112:2 117:9,11 129:3 135:10 141:21 145:10 148:25 149:10,16 155:17 156:19 <b>backed</b> 17:13 <b>backing</b> 126:13 <b>backwards</b> 143:3 <b>bad</b> 42:16 <b>balances</b> 28:24,25 29:9	<b>base</b> 136:6 <b>based</b> 39:9 51:3 53:14 54:2,10 56:13 58:21,23,23 61:23 64:13 66:2 87:10,15,15,18,18 87:22 89:10,23 90:23 94:25 102:1 102:21 146:22 147:24 148:17 149:8 152:18 <b>basically</b> 16:2 34:12 37:24 42:2 58:21 64:22 66:1 70:15 75:2,20 80:9 81:8 84:8 103:24 104:3 114:7 130:1 142:10 <b>basis</b> 5:6 20:21 28:8 52:5 55:9 61:22 62:2 75:6 75:22 81:2 133:1 <b>beef</b> 116:10 <b>beginning</b> 100:24 104:4 <b>behavior</b> 58:10 <b>believe</b> 8:25 18:24 24:2 25:8 26:5,21 26:23 35:1 44:5 56:9 59:12,20 74:18 83:21 89:20 89:20,22 149:10 <b>belongs</b> 22:17 56:20 <b>ben</b> 9:21 116:21 <b>ben's</b> 45:4 <b>best</b> 8:12,23 42:23 52:19 64:8 85:17 85:25 151:1
--	--	---	--

[better - certainly]

<b>better</b> 15:21 19:21 19:24 37:25 44:24 70:19 100:25 129:4 144:4 <b>beyond</b> 18:23 59:5 <b>bfalaw.com</b> 3:7,7 <b>big</b> 39:15 91:11 106:7 133:22,24 141:11,11 142:19 <b>bigger</b> 15:13 <b>biggest</b> 42:6 <b>binary</b> 127:2,6 <b>bit</b> 14:3 15:4 20:10 37:18 72:24 105:14 106:1 107:14,25 109:6 115:14 134:15 <b>blah</b> 145:18,18,18 <b>bleichmar</b> 3:4 <b>blob</b> 25:17 120:22 <b>blobs</b> 127:11,15 <b>blocks</b> 30:15 <b>board</b> 48:22 <b>book</b> 115:4 131:9 <b>books</b> 131:9,10 <b>bot</b> 26:5 <b>bother</b> 99:2 <b>bottom</b> 35:18 151:22 <b>boy</b> 141:11,11 142:19 <b>break</b> 5:10 60:13 60:19,23 63:8 93:12 117:5,7 130:17 133:8 <b>breakout</b> 6:19 <b>breaks</b> 5:5 <b>bring</b> 135:10 140:2 <b>broad</b> 14:21 20:19 88:9 89:6 138:10	141:7,9 <b>broadly</b> 82:22 88:8 103:15 119:18 131:5 <b>broke</b> 149:11 <b>broken</b> 72:23 <b>brought</b> 75:2,3 <b>browser</b> 89:1 155:11,12 <b>bucket</b> 49:4 58:20 <b>build</b> 20:22 33:24 34:13 41:7 83:16 91:9 108:11 138:24 139:14 140:13 141:24 144:4 146:20 <b>building</b> 28:7 110:3 143:13,17 <b>builds</b> 32:20 95:23 138:25 <b>built</b> 21:3 23:19 29:10 33:18,19,25 34:3 54:10 55:12 65:15 111:12 121:20 123:11 <b>bulk</b> 136:5 <b>bunch</b> 134:4 <b>burning</b> 7:13 <b>business</b> 64:21,23 65:5,23 66:6,6,21 66:22,23,23 67:15 68:3,17 69:3,5,13 70:3,5,20,20 71:6 71:7,11,13,15,25 72:12 73:11,12 75:4,7,7,18,20 76:3,16 77:4,8 78:19 79:5,11 80:12 81:10,13 82:5 83:16 94:23 114:12 139:10,16	139:20 140:21 142:2 149:17 <b>busy</b> 157:10 <b>button</b> 68:9,11 <b>buttons</b> 59:20 <b>buy</b> 81:16	<b>cares</b> 140:1 <b>cari</b> 3:10 <b>cart</b> 46:23 94:7,9 95:12,22 97:8,21 <b>case</b> 1:6 2:6 30:21 33:13 90:25 106:23 107:23 110:5 114:14 117:15 128:16 129:13,22 137:21 146:18,19,20,22 146:25 <b>cases</b> 16:20 37:10 37:21,21 51:14,15 53:6 103:19 136:3 <b>categories</b> 101:14 <b>categorizing</b> 12:25 <b>category</b> 42:19 48:14 49:1 60:15 60:15,24,25 61:1 64:5,6 86:4,20 89:14 93:7 <b>caucus</b> 93:9,13 <b>caveat</b> 108:6 <b>census</b> 146:2,4,9 146:10 147:13,14 148:4 151:23 152:3,7 <b>center</b> 128:15 129:3 <b>centers</b> 128:18,20 <b>certain</b> 11:1 37:21 45:13 48:6 67:9 68:3 79:17 102:1 108:12 127:6 141:14 151:19 154:3 <b>certainly</b> 7:9 45:21 47:1 154:17 159:7
--	--	--	--

## [certification - come]

<b>certification</b> 160:1	<b>chronologically</b>	58:7,23 59:4 61:8	110:22 114:15
<b>certified</b> 2:11 9:7	18:9	61:10,18,19 62:18	144:13 147:12
9:15,23 12:8	<b>chump</b> 135:3	62:22 64:16 71:10	<b>cleared</b> 126:24
160:4	<b>chunks</b> 127:5	71:24 72:5,15	<b>click</b> 15:24 16:15
<b>certify</b> 28:8 160:6	<b>circle</b> 93:9	73:9,12,18,23 74:3	94:18
160:15	<b>circumference</b>	74:17 79:2,23	<b>clicking</b> 88:19
<b>cetera</b> 8:16 115:19	45:9	83:2,13 84:9,22	<b>clicks</b> 49:9 50:15
148:5	<b>cite</b> 66:16	85:14,25 86:6,18	55:16,19
<b>chains</b> 101:25	<b>city</b> 90:20	87:4 89:12 90:4,7	<b>clients</b> 141:1
102:8	<b>clarification</b> 80:15	94:11,15 96:19	<b>close</b> 63:19 128:23
<b>challenge</b> 115:8	<b>clarifications</b>	99:25 100:12	<b>code</b> 10:22 28:10
117:18 131:4,20	12:22 86:22	103:18,25 104:6	40:20 72:9,10
<b>challenges</b> 120:12	<b>clarify</b> 16:11	107:5 108:6,19,22	100:7 117:16
128:17 130:9	17:22,23 37:9	110:3 112:15,21	129:14,21 148:17
<b>chance</b> 14:5,6	38:14 39:5 40:1	112:23 113:7	151:24
<b>change</b> 135:3	55:1 83:2 84:23	115:2,4,13 116:17	<b>cold</b> 43:22 132:4
<b>changes</b> 55:14	84:24 94:1 98:18	118:22 120:16,23	132:11,25 134:5,6
67:6 73:1 109:15	98:21 105:1 107:6	121:16,25 122:4	134:12,13,21
126:19 140:22	130:18,23	123:11 130:24	<b>collect</b> 81:9
<b>characters</b> 106:1	<b>clarifying</b> 6:17	131:18,25 132:10	<b>collected</b> 13:13
<b>cheap</b> 91:3	12:17 43:7 115:13	132:15,19 133:16	49:4 50:9,20 51:6
<b>cheaper</b> 38:2	<b>clarity</b> 27:18	137:17 140:18	53:25 69:3,12
<b>check</b> 31:1,8 45:1	44:16	141:3,25 142:25	70:22 71:9 74:13
91:24 124:23	<b>clark</b> 3:16 5:24,25	144:18 145:11	82:5,8,10,12 144:2
<b>checkout</b> 94:19	8:15 10:1,6 11:23	146:1 149:10	153:22
95:4 149:21	12:1,4,6 14:24,25	150:17,22 151:11	<b>collecting</b> 49:25
<b>checkpoint</b> 155:14	15:11,14,19 16:4	156:11	50:4 81:8
<b>checkpoints</b>	16:13 17:8,20	<b>class</b> 23:8 143:5	<b>collects</b> 47:2 80:13
140:24	20:16 21:9,23	<b>classifications</b>	<b>color</b> 45:8 46:6
<b>checks</b> 20:25	22:8,12 23:13,24	26:3	<b>color's</b> 45:6
28:24,25 29:9	24:3,17 26:1 27:6	<b>classifiers</b> 23:15	<b>column</b> 43:10,12
125:11	27:10,18,24 28:6	<b>classify</b> 26:6 86:1	45:15,16 112:13
<b>chen</b> 4:2	28:25 29:24 30:11	135:22	114:18 147:22,22
<b>chhabria</b> 158:22	31:2,4,12,17,21	<b>claufenberg</b> 3:13	<b>columns</b> 43:9,15
<b>choose</b> 36:1 37:19	33:8 41:16 42:1,4	<b>clear</b> 10:11 11:15	112:8 114:24
97:20 120:7	42:23 43:3,6 44:5	13:9 38:20 40:4	<b>combination</b> 51:9
<b>chooses</b> 129:20	44:10,18,23 47:10	40:24,25 41:1	56:2 89:24
<b>chose</b> 46:7 53:19	49:6 50:12 51:9	42:22 47:20 60:24	<b>come</b> 29:21 42:12
69:11	51:11 54:4 55:1	77:21 78:5 80:11	53:22 98:2 120:21
<b>chosen</b> 49:15	55:10 56:1,13	83:15 104:18	124:19 146:18
64:23	57:1,6,12,15 58:2	105:4 107:8,13	



[comes - cookies]

<b>comes</b> 14:12 46:16 54:1 83:1 91:24 119:3 141:10 144:18 149:15 <b>comfort</b> 10:15 19:12 <b>comfortable</b> 8:11 11:12 19:17 <b>coming</b> 53:11 129:3 <b>command</b> 80:18 <b>commencing</b> 2:12 <b>comment</b> 10:12 29:14 123:1,2,6 125:12,13,15,16 <b>comments</b> 87:22 123:3,4,5,6 <b>commingle</b> 136:7 <b>commit</b> 31:1,8 129:2 <b>commitment</b> 112:16 <b>commodity</b> 76:19 <b>common</b> 34:1 129:22 142:12 <b>commonly</b> 141:13 143:20 <b>communicate</b> 158:23 <b>communication</b> 89:4 <b>community</b> 65:11 65:14 <b>companies</b> 66:15 <b>company</b> 33:16 96:7 149:18 <b>competency</b> 95:24 <b>complete</b> 13:5 15:7 31:23 44:12 58:17 85:18 119:8	<b>completed</b> 28:12 28:12 <b>completely</b> 10:20 10:23 11:3 65:19 72:4,5 82:16 137:25 <b>completion</b> 29:2 <b>complex</b> 6:24 <b>compliance</b> 42:13 <b>complicated</b> 42:18 111:13 <b>components</b> 74:19 <b>comprehensive</b> 119:12 <b>compress</b> 134:14 <b>compressed</b> 61:25 132:4 <b>computation</b> 38:3 135:23,24 136:1,4 136:5,9,10 <b>computational</b> 131:16,19 <b>computations</b> 37:12,13 136:12 136:15 <b>compute</b> 37:16 <b>computed</b> 87:2 89:10 <b>computer</b> 65:21 <b>computing</b> 37:15 93:3 102:9 <b>concept</b> 23:4 57:15 <b>conceptually</b> 123:21 <b>concrete</b> 23:13 49:7 87:6 122:3 137:20 <b>confidential</b> 1:9 158:4,10	<b>configuration</b> 128:4 <b>confirmation</b> 155:15 <b>conflict</b> 125:11 <b>confused</b> 63:5 137:22 <b>confusion</b> 111:8 <b>connect</b> 70:17 <b>connection</b> 130:25 <b>connects</b> 89:12 <b>consent</b> 73:3 <b>consider</b> 93:13 <b>consistent</b> 28:8 54:7 <b>constant</b> 103:23 <b>constantly</b> 13:20 55:13 <b>constraints</b> 134:4 142:5 <b>construct</b> 42:11 <b>consult</b> 28:21 <b>consulted</b> 22:6 <b>consumed</b> 36:3 <b>consumer</b> 1:6 2:6 79:21 83:4 <b>contain</b> 114:17 116:1 127:13 <b>contains</b> 15:10 45:15 78:1 112:12 145:4 <b>content</b> 16:1,19,23 17:3,24 24:25 53:5 82:24 93:25 94:2 95:11,15,18 95:22 96:2,14,18 96:24 97:2,7,7 98:1,3,15 122:24 122:25 124:4,5,7,9 124:18 149:19	<b>context</b> 21:25 <b>contexts</b> 151:19 <b>continent</b> 129:2 <b>continually</b> 55:11 <b>continue</b> 29:7 66:13 67:20 155:16 <b>continuing</b> 154:14 <b>contract</b> 39:12 <b>contracts</b> 68:1 139:24 142:3 <b>contrast</b> 135:19 <b>control</b> 19:6 95:6 114:7 148:7 <b>controls</b> 54:17 123:13 142:6 <b>convenient</b> 117:5 <b>conver</b> 77:8 <b>conversation</b> 131:23 132:3 135:9 148:25 <b>conversations</b> 6:20 154:17 <b>conversion</b> 64:25 70:2,4 74:10,24 77:8 78:11 80:7 80:12 81:6,8 82:2 82:11,21 115:24 <b>conversions</b> 68:18 78:13,24 79:7,7,14 114:14 <b>convert</b> 75:21 114:4 <b>convey</b> 158:25 <b>cookie</b> 21:18,18 150:12 155:6,7,14 155:17,19 156:2 <b>cookies</b> 47:21 114:13 115:24,25 116:1,1,3 150:11 154:18,19,23,23
--	--	---	---

## [cookies - data]

154:25 155:3,3,23 156:3,4,15,15 <b>coral</b> 80:18 97:15 <b>core</b> 106:6 <b>corley</b> 13:12 <b>corned</b> 116:10 <b>correct</b> 15:14 16:4 18:14 22:8 25:19 26:24 27:10,24 31:2,20,21 34:15 35:20 42:1 47:25 56:1 60:8 62:22 62:23,24 72:6 73:18 76:17 84:9 85:25 86:6,18 98:8 107:17 113:7 120:23 126:2 138:4 158:9 <b>correction</b> 34:18 <b>correctly</b> 30:2 <b>costco</b> 96:20 <b>counsel</b> 3:1 6:13 6:19 7:12,13,19 10:19 11:9 13:1 46:18 47:18 48:11 49:23 60:17,20 61:3 110:25 115:20 116:5 117:1 135:9 <b>couple</b> 5:15 15:22 18:18 22:12 52:21 87:6 90:7 103:3 131:9 153:2,5 <b>course</b> 117:6 142:17 <b>court</b> 1:1 2:1 5:5 5:17,18 8:7,9,11 8:13,18 9:2,10,18 10:1 12:3 13:11 29:11 63:20 76:10 92:15 117:4	157:22 158:12 160:1 <b>court's</b> 6:9 <b>cover</b> 54:2 156:22 <b>covered</b> 154:16,16 <b>covers</b> 154:11 <b>crc</b> 1:18 160:25 <b>create</b> 28:9 41:12 41:17 79:9 109:6 119:1,6 120:6 129:4 <b>created</b> 29:2 38:19 101:25 123:25 124:8 148:3 <b>criteria</b> 21:12 59:5 147:7 <b>cross</b> 27:20 <b>crr</b> 1:18 160:25 <b>crutcher</b> 3:16,20 <b>csr</b> 1:18 160:25 <b>cue</b> 15:20 <b>cued</b> 121:7,8 <b>cup</b> 141:15 <b>curiosity</b> 86:25 <b>currency</b> 98:3 <b>current</b> 35:12 <b>currently</b> 24:9 33:13 89:20 <b>custom</b> 79:9 82:25 83:5,17 84:18,25 85:3,12,18,23,23 86:4,16 95:5,6 97:6,18,24 99:3,10 99:11 100:10 129:14 149:16 <b>customer</b> 139:5 140:9 141:9 <b>customers</b> 138:23 139:1 <b>customized</b> 140:8	<b>cut</b> 15:7 150:1 153:10 <b>cycle</b> 74:17,21 146:17 <b>cycles</b> 93:3 <b>d</b> <b>daily</b> 133:1 <b>daniel</b> 1:11 6:2 <b>dapl</b> 85:1 <b>data</b> 6:6,8,9,10 7:6 13:3,13,13,14,16 14:7,9,12 15:7,8 15:10,17 16:3,8,22 17:6 18:2,3,10,21 18:22 20:8 21:1 21:12,13,20 22:7 22:10,14,18,25 23:20 24:19 26:17 27:4,15 28:3,21 30:19,23 32:18,19 32:20,21,21,24,25 33:2,5,6,7,10,17 33:19,21,23,23,24 33:24 34:3 35:3 35:10,11,22,24 36:4 37:6,11,12,23 39:7,18 40:17,18 41:7,18,18,19 42:2 42:7,8,14,20,21,24 44:13,20 47:7,8,18 47:23 49:3,4,11,11 49:12 50:2,19,23 50:24 51:1,2,5,15 52:3,3 53:7,12,16 53:20,21,25 54:8,9 54:18,21,24 55:3,5 55:24 56:18 57:5 57:7,7,10,19 58:6 58:19,20 59:9,10 60:4 64:6 69:3,11 69:12,14,19,20,23	70:22 71:9,17 73:2,5 74:21,24 75:1,15,19 76:5,8 76:18,18 77:4,11 78:4 80:14,16 81:8,9,19 82:3,4,5 82:8,9,10,12,13 83:11 86:13 87:1 87:1,25 88:2,3,5,7 88:8,10 89:8,9,15 89:18,18,25 90:3 90:11,17,22,23 91:7,20 92:11,19 93:2,7,25 94:2 95:11,18 96:8,9,12 96:14 98:15 99:6 99:14,15 100:10 101:22 102:19,24 103:2,5,7,19,22 106:16 108:9,12 108:15,16 110:10 110:15,15,18 111:10 112:2,5 113:15 118:16 119:4,10,16,17,17 120:13,14,20 121:20,22 122:19 123:22 124:21 125:2,9,21,21 126:15 127:2,12 127:13 128:15,18 128:19 129:3 130:10,14,20 134:21,22 135:3,4 135:15 136:25 137:14 142:6 143:13,21 144:2 144:14 145:4,7,9 145:10 146:3,4,5 146:10,10,10,11 146:14 147:10,13
---	---	---	--

[data - developed]

147:15,15 148:4 148:24 149:3,5,24 150:4 151:12,23 151:25 152:15,16 153:9,21 154:22 <b>database</b> 52:12 66:19 109:20 120:9,9 125:9 128:7 <b>databases</b> 51:16 69:18 81:24 105:15 <b>datas</b> 44:2 61:24 <b>dataset</b> 36:3,15,15 38:8 39:7 76:3 130:3,5 131:7 146:21 <b>dataset's</b> 37:24 <b>datasets</b> 88:1 128:13 145:16,17 147:14 <b>dataswarm</b> 35:6 35:19 41:13 75:14 101:2,8,15 102:12 102:19 129:11,12 <b>date</b> 94:20 117:23 143:1 149:15 160:18,19 <b>datr</b> 154:23 155:6 155:7,14,17 <b>david</b> 4:8 <b>davis</b> 3:5 <b>day</b> 38:5 131:14 131:15 134:6,9 136:5,12 153:17 157:10,14 <b>day's</b> 145:7 <b>days</b> 15:22 67:9 112:15 <b>de</b> 56:21	<b>deal</b> 67:11 101:8 <b>deals</b> 139:18 <b>decade</b> 109:14 <b>decide</b> 91:23 <b>decided</b> 40:8 <b>decides</b> 52:16 146:19 <b>decision</b> 27:13 28:2 29:21 90:23 96:13,13,16 156:3 <b>decisions</b> 29:19,20 60:9 <b>declaration</b> 45:13 <b>decode</b> 116:2 <b>dedicated</b> 18:13 127:11 <b>deem</b> 120:18 <b>default</b> 87:12 <b>defendant</b> 3:15 4:1 <b>defer</b> 17:22 99:22 <b>define</b> 17:15 23:11 36:15 43:15 69:22 <b>defined</b> 47:17 48:25 57:1,4 88:8 108:14 123:7 151:18 <b>defines</b> 55:23 56:17 <b>defining</b> 24:11 33:7 47:8 90:8 <b>definitely</b> 13:4 121:16 <b>definition</b> 16:10 19:5 46:20 47:7 47:11,15,24 48:3 88:2 109:4 <b>definitions</b> 117:17 <b>delete</b> 16:20 123:2 123:3,3 125:13,15 125:16	<b>deleted</b> 110:10,13 110:19 111:10,22 112:1 122:25 <b>deleting</b> 134:22 <b>deletion</b> 112:16,20 113:3 122:24 <b>deliberately</b> 45:14 <b>delineate</b> 88:10 <b>delineated</b> 60:2 88:12 <b>delineation</b> 144:13 <b>delivered</b> 134:25 <b>demand</b> 136:4 <b>denied</b> 29:22,24 <b>denver</b> 89:22 <b>departed</b> 49:2 <b>depend</b> 69:17 121:21 <b>dependency</b> 101:24 102:7,8 <b>depending</b> 25:5 43:15 69:9 87:5 92:18 121:9 145:6 <b>depends</b> 24:15 26:17 51:12 53:11 56:16 88:22 90:21 111:14 114:12 <b>deployed</b> 29:3 <b>deprecate</b> 55:19 56:23,24 <b>deprecated</b> 35:1 55:12 61:23 62:16 65:6,7,18 66:22 67:5,10 <b>deprecation</b> 67:7 <b>depth</b> 128:1 <b>derek</b> 3:10 <b>derivation</b> 146:9 <b>derivative</b> 25:2 <b>derive</b> 87:11,15,19 87:20,23 88:20	147:15 <b>derived</b> 87:4,5 88:1,2,7,8,10,14 89:9,25 90:3,11,17 91:7,20 146:3 147:12,16 151:23 151:25 152:19 <b>deriving</b> 146:10 <b>describe</b> 35:7 121:12 146:14 148:2,2 <b>described</b> 48:12 94:4 <b>describes</b> 14:7 <b>describing</b> 45:10 <b>description</b> 143:15 <b>descriptions</b> 118:15 <b>descriptively</b> 61:6 <b>design</b> 43:17 108:10 <b>designated</b> 158:10 <b>designating</b> 158:4 <b>designed</b> 138:6 <b>desire</b> 7:13 159:3 <b>detail</b> 7:4 71:17 <b>details</b> 28:10 79:9 122:7 126:14 <b>determination</b> 21:11,22 40:13 122:20 <b>determinations</b> 21:8 <b>determine</b> 20:12 21:21 60:3 103:20 114:17 152:3 <b>determined</b> 47:2 55:22 <b>developed</b> 21:4 72:18,23
---	--	--	--



**[developer - dynamic]**

<b>developer</b> 28:20 34:21 65:23 66:15 67:19 68:4 71:16 72:1 73:24 95:23 96:9 122:13 123:22 125:23 127:22 129:20 140:20 <b>developers</b> 36:1 38:22,23,24 39:3 64:16,17,18,21,22 95:24,25 100:5,7,8 114:24 125:19 <b>development</b> 33:20 65:14 140:10 142:2,9,13 <b>deviations</b> 142:9 <b>device</b> 24:7 49:11 89:1 114:13 <b>devices</b> 54:24 56:14,25 <b>devoted</b> 32:4 <b>diagnose</b> 41:21 <b>difference</b> 114:5 <b>different</b> 13:2 14:23 15:6 22:18 24:8,10 47:8 48:23 51:1,1,2 52:14 69:10 72:3 72:4,5,8,11 78:10 79:16 85:1,22 96:11 98:22 118:15 122:1 126:3 132:3 133:6 135:20 138:12 145:18 150:12 154:8 155:21 <b>differentiate</b> 18:16 43:21 <b>differentiated</b> 46:12	<b>digital</b> 70:17 81:12 <b>direct</b> 8:24 137:1 <b>direction</b> 40:13 160:13 <b>directly</b> 47:5,11 54:13 57:12 58:25 103:5 124:21 126:21 137:6 144:18 149:19 <b>discard</b> 100:11 <b>disclosed</b> 158:6 <b>discoverable</b> 6:8 13:12 <b>discovery</b> 115:9 <b>discuss</b> 6:4 45:20 <b>discussed</b> 11:16 32:10 77:24 142:16 <b>discussion</b> 29:16 32:2 142:21 <b>disk</b> 29:3 154:5 <b>disks</b> 30:15 76:20 <b>dismiss</b> 157:22 <b>displayed</b> 81:13 <b>distinct</b> 133:9 <b>distinction</b> 18:18 77:22 80:19 99:17 <b>distinguishing</b> 71:1,5 <b>distributed</b> 17:13 84:5 106:4 107:15 108:25 109:2,20 120:9 126:17 <b>district</b> 1:1,2 2:1,2 <b>dloeser</b> 3:12 <b>dna</b> 24:24 26:13 26:15 <b>document</b> 27:2,23 83:22 <b>documentation</b> 70:21 71:7,8,15,19	72:2,16 77:3,7 79:10 80:4,4 83:10,15,18 94:22 97:13 101:24 117:13,18,19,22 117:24 118:3,3 129:6 133:6 149:9 <b>documented</b> 73:17 97:11,12 141:19 141:24 <b>documenting</b> 117:16 <b>documents</b> 70:18 <b>doing</b> 8:7 39:8 40:13 93:1 99:2 106:18 143:22 153:23 154:2 <b>domain</b> 132:6 <b>double</b> 45:1 <b>download</b> 15:2 46:25 83:4 119:21 <b>downloadable</b> 16:5 <b>downloaded</b> 52:1 <b>downloading</b> 15:4 <b>dozens</b> 32:24 <b>draconian</b> 92:11 <b>drafts</b> 61:6,6 <b>draws</b> 149:25 150:5 <b>drift</b> 117:19 137:11 <b>drink</b> 63:13 <b>drive</b> 77:1 79:10 <b>driven</b> 56:21 139:5 <b>drives</b> 134:20 <b>drop</b> 25:17 <b>duly</b> 9:6,14,22 12:7	<b>dump</b> 36:12 <b>dumped</b> 36:14 <b>dumping</b> 118:4 <b>dunn</b> 3:16,20 <b>duplicated</b> 105:22 <b>duplicative</b> 21:18 21:18 22:7 59:24 106:17 127:19 <b>dyl</b> 13:25 14:8,9 14:23 15:1,5,8,10 15:13,19 18:23 20:8,13 21:6,20 22:5,10 24:22,22 25:7,11,24,24 26:2 27:4,15 28:22 30:22,24 31:10,14 31:22,23,23,24 32:3,8 38:20 42:21 43:1,5 44:4 44:14 45:17 46:7 48:19 51:22 52:1 56:5,20 57:20,21 57:21 58:6,20 60:5 64:13 73:6 75:3,10 77:12 78:3 79:25 83:4 84:20 85:12,24 86:2,17 89:22 93:25 94:10,12 95:18 96:16 97:3 98:6,15 99:5,13 103:17 120:13,20 147:5,8,18,25 148:4,14 149:3,8 151:12,23 152:11 152:16,20 154:6 155:22 156:2,4 <b>dynamic</b> 55:11
---	--	---	--

[earlier - executed]

<b>e</b>	<b>ends</b> 53:9 54:14 84:19 90:20 91:8 92:13 130:7	<b>entry</b> 44:3	<b>exactly</b> 38:4 49:7 49:8 53:18 68:2 77:1 81:15 86:10 92:4 107:22 111:2 118:6 124:2 128:25 139:7,18 148:20
<b>earlier</b> 41:20 69:5 91:10 92:7 94:8 131:8	<b>enforce</b> 122:9 124:15,16	<b>ephemeral</b> 23:4 136:22	<b>examined</b> 9:7,15 9:23 12:8
<b>easier</b> 87:7 159:4	<b>engaged</b> 156:12	<b>equal</b> 124:8	<b>example</b> 22:14 23:13 24:13,24 26:7 35:9 37:15 41:22,23 45:11 46:21 52:23 54:22 56:6 58:9 59:15 59:23 64:24 66:2 68:8 72:19 81:23 84:4 91:1,10 94:8 96:15 104:1 106:17 120:19 122:3 123:1,10,24 124:6 137:18,20 141:2,4 149:5
<b>easy</b> 34:8 37:14 68:22	<b>engineer</b> 20:14,20 21:4,7,9,17,21 22:10 23:5,19 27:3,9 28:18,21 31:1 38:22 39:8 39:13,14 41:6 44:6 54:16 70:5 71:16 83:25 84:16 110:3 116:16 122:3,6 145:11 146:14	<b>equals</b> 124:17	<b>examples</b> 22:13 49:7,20 50:6,12,18 51:23 52:21 53:23 54:2 87:6 89:6 90:5 95:8,16 101:2,5 123:20
<b>echo</b> 150:3	<b>engineer's</b> 44:11	<b>equivalent</b> 38:1	<b>excel</b> 80:19
<b>economical</b> 115:17	<b>engineering</b> 7:8 8:22 14:16,18 20:7 22:5 25:21 41:11 45:2 48:3 56:19 71:21 123:10 143:10	<b>error</b> 125:18	<b>exception</b> 142:18
<b>educate</b> 157:7	<b>engineers</b> 6:13 7:18 10:20 11:19 11:21 20:12 28:22 34:3,4 42:11 61:11 99:22 101:7 101:13,15 108:10 120:11 150:23 157:6,9	<b>especially</b> 26:12 84:9,14 108:3 110:16	<b>exceptions</b> 142:8
<b>effective</b> 18:8		<b>esq</b> 3:4,5,10,10,16 3:21 4:2,3	<b>excess</b> 26:22
<b>effectively</b> 16:3 26:14 33:11 35:7 35:19 38:11 53:15 75:2,9 77:9,16 81:16 120:8 136:19		<b>essence</b> 43:13	<b>exclude</b> 14:14
<b>efficient</b> 64:11 127:5 131:3		<b>essentially</b> 120:18	<b>excluded</b> 45:22 59:3,8,14 60:5
<b>effort</b> 7:8 74:4 158:24		<b>et</b> 8:16 115:19 148:5	<b>excluding</b> 43:10
<b>either</b> 28:13 30:6 74:4 79:8 90:18 92:12 116:18,19 116:23		<b>eugene</b> 9:5 52:8 67:3 107:24 108:1 122:21 131:1	<b>excuse</b> 147:19
<b>elements</b> 28:11 30:1 74:20 123:13 149:4,5		<b>evaluations</b> 20:25	<b>executed</b> 35:4,5 114:8
<b>else's</b> 103:22		<b>event</b> 32:19 33:23 79:8,10 86:5 94:19 95:3,3 97:6 97:6,15,19,24 100:10 103:23 149:15	
<b>email</b> 3:7,12,18,23 4:4 36:6 84:10 104:11,23 155:15		<b>events</b> 49:10 50:18 79:18 95:21 96:21 97:25 99:10,11	
<b>emailing</b> 36:5		<b>eventually</b> 52:1,9 75:10 81:2 103:5 121:6 146:15	
<b>embed</b> 26:17	<b>engrossed</b> 5:8	<b>everybody</b> 5:16 10:24 24:10 29:5 31:7 63:8,16,18 93:17 117:2 143:10 157:5,22	
<b>embedded</b> 25:3 118:1	<b>ent</b> 121:13,18,21 122:2,8,10,16,19 123:8 124:19,24 125:1,3	<b>everybody's</b> 5:14 153:18	
<b>employee</b> 20:19 29:6,7 42:7 160:16	<b>entire</b> 21:19 59:8 112:22 130:8,12 130:13 134:6 136:6,6,10	<b>exact</b> 30:14 76:22 142:25	
<b>ended</b> 159:11			

## [exhibit - financially]

<b>exhibit</b> 14:15 35:17 145:15 150:14,18,20 151:9 152:8 <b>exist</b> 73:15 74:6 83:19 146:19 156:16 <b>existed</b> 62:21 <b>existing</b> 39:17 41:19 122:14,18 <b>exists</b> 21:13 22:22 29:4 62:20 71:3 111:15,24 112:1 112:11,12 <b>expansive</b> 91:1 <b>expect</b> 10:24 11:4 <b>expectation</b> 11:18 <b>expected</b> 10:16 <b>expedite</b> 7:10 <b>experience</b> 59:25 90:25 91:10,12,15 129:4 143:18 144:21,24 <b>experiences</b> 59:13 143:14,17 144:4 <b>expert</b> 88:25 126:14 141:5 <b>expertise</b> 77:14 128:5 132:7,8 <b>experts</b> 46:25 <b>explain</b> 14:20,22 96:4 98:25 99:24 122:5 137:24 150:15 153:8 <b>explained</b> 77:10 81:3 96:10 132:18 <b>explaining</b> 122:4 153:21 <b>explains</b> 77:9 83:22	<b>explanation</b> 126:24 <b>explicit</b> 116:24 <b>export</b> 14:8 37:18 <b>extensive</b> 37:11,16 37:23 71:14 128:20 129:23 135:23 136:12 <b>extensively</b> 129:1 <b>extent</b> 100:15 <b>external</b> 69:6 138:8 <b>externally</b> 137:10 <b>extract</b> 106:25 <b>extracted</b> 14:9 51:25 52:9 69:13 69:23 74:14 82:13 107:2 <b>extraction</b> 106:9 108:8 <b>extremely</b> 157:16 <b>eye</b> 45:6,8,9,10 46:6 <b>eyes</b> 1:10 133:23 <b>f</b> <b>facebook</b> 1:6 2:6 4:2 6:14,23 7:7 8:21 15:21 16:1 16:17 25:1 26:24 29:6,7 32:20 34:12,16,19 38:22 39:8,14,21 46:22 47:2,13,22,23 48:15 52:22 62:4 62:7 64:24 65:23 66:5 67:15 68:10 68:23 70:20 74:11 77:3 80:16 83:11 84:5,11 88:15 90:9,13 99:2,3,5,6 99:14,20,21 102:4	102:11 103:16 105:4,5 109:4,6 110:1 114:1,2,20 116:1,2 119:9 122:12 125:17 135:2 136:6 138:25 141:17 145:17 149:25 150:5 153:7 155:2 157:6,9 <b>facebook's</b> 32:20 33:23 62:4 65:5 <b>facing</b> 69:6 <b>fact</b> 16:14 52:24 53:4 <b>facto</b> 56:22 <b>factor</b> 90:3 <b>factors</b> 96:10 <b>factual</b> 151:12,20 <b>fair</b> 46:1 91:5 <b>fairly</b> 132:25 <b>familiar</b> 21:19 132:25 <b>fancy</b> 25:5 <b>far</b> 32:10 128:20 142:4 <b>fast</b> 92:5 103:13 129:24 136:1 <b>fastest</b> 61:15 <b>fb.com</b> 4:4 <b>feasibility</b> 115:6,8 <b>features</b> 68:3,6 97:23 139:1 <b>federal</b> 160:2 <b>feed</b> 49:16 <b>feel</b> 11:12 19:17 38:17 <b>feeling</b> 20:3 151:1 <b>field</b> 30:19 43:23 44:1,21 45:9,23,23	<b>fields</b> 14:13,13 43:1,4,19 59:7 60:3 83:19 95:1,5 95:7 118:15,16 149:14,16 <b>figure</b> 6:7 24:18 25:21,23 29:20,22 54:21 57:19 60:6 64:12 65:4 74:9 80:13 82:15 83:1 85:10 87:2 89:8 101:19 119:15 138:21 140:3,5,6 141:18 143:14 148:13 154:3 156:14 <b>figured</b> 90:15 156:20 <b>figuring</b> 106:20 146:17 <b>file</b> 13:25 14:9 15:2,5 22:4,7,11 25:7 30:24 36:12 43:2,5 44:4 48:19 56:14,20 57:20 58:20 60:5 64:13 84:20 85:24 86:2 86:17 94:10,12 96:17 98:7 99:5 99:13 103:17 120:13 130:12 147:5,18 148:4,14 149:3,8 154:6 155:22 156:2,4 <b>files</b> 23:8 <b>filter</b> 96:19 <b>filtered</b> 14:12 <b>filtering</b> 96:21,25 <b>finally</b> 90:15 <b>financially</b> 160:15
--	---	--	--

[find - garrie]

<b>find</b> 23:3 53:15 115:4 130:20 131:11 <b>finding</b> 131:7 <b>fine</b> 58:4 62:8 82:20 <b>finish</b> 13:4,7 122:4 135:11 143:19 <b>finished</b> 64:5 <b>fire</b> 97:14 <b>fired</b> 108:20 <b>first</b> 6:21 9:6,14 9:22 12:7 13:24 13:24 20:18 47:12 47:21,23 48:14 49:15 59:17 60:19 85:11 88:14 89:16 89:24 91:12,19 101:11 104:10,22 106:13 116:6 117:13 134:22 136:24 137:13 138:3,6 155:13,19 <b>five</b> 39:18 63:8,11 66:4,6 117:12 131:13,14,15 <b>flagging</b> 58:12,14 58:16 <b>flavor</b> 24:10 <b>flips</b> 128:13 <b>flow</b> 60:16 <b>flows</b> 152:4 <b>focus</b> 116:13 <b>focuses</b> 77:20 101:15 139:9 <b>folks</b> 72:21 156:12 <b>follow</b> 29:9 32:17 141:25 142:1,5 147:3 154:13 156:21,21	<b>followed</b> 139:25 140:17 141:16,23 <b>following</b> 157:4 <b>follows</b> 9:8,16,24 12:9 142:22,22 <b>fonti</b> 3:4 <b>foregoing</b> 160:7,9 160:14 <b>foreign</b> 105:17 <b>forensics</b> 26:18,24 <b>forget</b> 8:6 <b>forgetting</b> 109:1 <b>forgot</b> 106:7 130:17 <b>form</b> 76:3 <b>formalities</b> 12:12 <b>format</b> 6:12 8:2 16:5 34:20 <b>forth</b> 13:19,20 119:6 141:22 156:19 160:8 <b>forward</b> 28:17 82:23 154:14 <b>found</b> 31:10 146:1 <b>four</b> 8:23 103:14 103:16 104:7,18 157:9 <b>fraction</b> 33:1 <b>frame</b> 125:7 <b>frames</b> 23:23 24:7 90:10 <b>framework</b> 22:3 <b>frameworks</b> 121:18,19 122:9 122:15,18 <b>francisco</b> 3:22 <b>frankly</b> 138:16 143:8 <b>free</b> 8:20 10:12 68:9	<b>freeze</b> 58:1 <b>freezing</b> 58:3,3 <b>friends</b> 124:11,13 124:13 <b>front</b> 69:7 127:21 <b>froze</b> 10:3 <b>fulfilled</b> 137:5 <b>fulfilling</b> 143:23 <b>full</b> 74:17 115:9 158:19 <b>fully</b> 13:6 35:1 48:21 67:5 103:20 135:6 158:21 <b>fulmer</b> 1:18 2:11 160:24 <b>function</b> 15:5 124:1,2,14 <b>functional</b> 27:20 <b>functionality</b> 36:10 108:17 <b>functioning</b> 109:13 <b>functions</b> 106:6 149:1 <b>funny</b> 29:13 <b>further</b> 6:4 12:22 30:23 61:20 86:22 93:7 115:22 154:17 160:13,15 <b>future</b> 35:2 90:22 92:23	19:21,24 20:2,5 21:7,15 22:2,9 23:2,20 24:1,4,18 25:3,9,16,20 26:16 26:20,23 27:8,11 27:22,25 28:18 29:11,13,17 30:8 30:25 31:3,6,13,25 33:5,17,21 34:10 34:23 35:3,14,21 36:8,16,22 37:1,3 37:5,24 38:7,12,13 38:24 39:2,6,22 40:5,8,12,16,23 41:4,24 42:2,10,25 43:4,12 44:9,17,21 44:25 45:18,24 46:2,8,11,16 47:4 47:14 48:1,5,16,21 49:22 50:22 52:2 53:14,21 54:9 55:3,18 56:4,16 57:4,9,14,17,24 58:4,19 59:2,6 60:2,9,12,17,22 61:9,21 62:6,10,15 62:25 63:4,15,21 63:24 64:3,20 65:2,8,11,13 66:1 66:20 67:14 68:13 69:2,12,22 70:1,15 71:12,20 72:3,7 73:7,10,16,19 74:1 74:7,23 75:13,23 76:8,11,13 77:2,7 77:15,19 78:6,10 78:15,18 79:1,22 80:3,20 81:4 82:1 82:8,20 83:10 84:2,17 85:3,9,20 86:10,19 87:25
		<b>g</b>	
		<b>gar</b> 100:6 <b>garbage</b> 100:7 <b>garrie</b> 1:11 5:4,17 5:19,21,24 6:1,2 8:5,14,17,19 10:4 10:11,18 11:11,17 12:11 15:8,12,15 16:2,6 17:4,9 18:4 18:12,16,25 19:9	

[garrie - graph]

88:6 89:7 90:1,5 90:14 91:5 92:2,7 92:16,17,25 93:6 93:10,11,16,21 94:3,13 95:10 96:5,7,23,24 97:11 98:6,9,17,24 99:1 99:18 100:2,18 101:6,16,19,21 103:23 104:3,7,16 104:21 105:7,23 106:7,14,19,22,24 107:10,14,17,19 108:18,20,25 109:16,19,23 110:7,24 111:4,7 111:20 112:7,14 112:18,22 113:5,8 113:19,22 114:2,9 114:16,22 115:10 115:15 116:20 117:1,4,11,24 118:5,10,17,23 119:2,12,15,22 120:13,24 121:7 121:12,24 122:1 123:9,18,19 124:16 125:1,24 126:4,8,23 127:12 127:18,23 128:6 128:12,22 129:5 129:25 130:16 131:12,22 132:1 132:13,18,21,22 133:3,22 134:1,8 134:11,24 136:16 136:23 137:7,12 137:24 138:15 139:12 140:3,25 141:6 142:10,20 143:3,19 144:1,8	145:2,14,21,24 146:2,9 147:3,9,20 148:10,18,23 149:23 150:4,9,19 150:25 151:6,9,14 151:16 152:1,7,13 153:2,13,25 154:9 155:2,20 156:1,13 157:8 158:8,14,21 159:6 <b>gates</b> 142:16 <b>general</b> 38:18 57:15 85:7 132:9 146:25 <b>generalized</b> 34:2 <b>generalizing</b> 134:1 <b>generally</b> 19:2 114:22 126:21 128:10 141:24 <b>generate</b> 26:14 105:16 <b>generated</b> 13:25 52:3 <b>generating</b> 19:15 30:24 41:18 82:24 <b>generation</b> 14:9 <b>geo</b> 88:10 <b>geographical</b> 128:14 <b>geographically</b> 128:22 <b>getting</b> 78:22 93:22 99:7 125:2 137:6 142:11 <b>gibberish</b> 99:15,18 99:23 <b>gibson</b> 3:16,20 <b>gibsondunn.com</b> 3:18,23 <b>gigs</b> 15:22	<b>give</b> 5:10 10:15 11:20 23:7 28:10 30:25 46:16 56:6 63:4 66:2 70:6 84:3 103:22 153:4 154:3,4 158:19 <b>given</b> 94:25 106:8 118:12 143:5 144:1 145:2 153:4 158:5 <b>gives</b> 22:3 <b>giving</b> 70:16 95:16 120:8 <b>globally</b> 126:17 128:18 <b>go</b> 5:4 6:14,19 8:25 11:5 12:16 13:4 20:22 26:2 29:15,17 30:23 35:12 36:23 37:3 37:4,5 44:20 52:23 60:6 62:1 63:13,16,21 64:3 67:2,3 83:3 91:16 93:18 94:17 96:16 100:21 101:4 103:13,15 104:14 109:25 110:1,2 117:7 118:8 119:1 121:10,10,11 122:13 124:9,24 128:3 131:4 133:20 140:23 142:14,15 144:22 146:20 149:10 153:6,17,19 159:8 <b>goes</b> 21:8 36:21,24 44:11 54:4,14 56:19 69:4 72:25 77:13 78:3 90:15 90:25 91:3 97:16	108:9 131:24 142:4 152:20 153:9 <b>going</b> 6:25 7:23 8:1,14,19 10:16,18 10:19 12:14,15,20 12:23 13:8,20 17:21 26:9 29:14 32:15 41:21 45:19 46:13 47:10 48:16 48:23 49:19 50:16 51:6,10 55:11,13 56:18 60:13,15,19 60:20,22 61:9 63:7 64:6 66:11 72:23 76:22 77:20 79:19 83:24 84:11 91:14 93:11,11,16 94:16,22 110:6 115:11 119:10 120:21 137:19 141:21 143:9 151:4 152:16 153:13,22 156:3 <b>good</b> 5:17,20,21 5:23,25 23:13 31:8,9 39:23 41:9 49:20 56:7 72:19 93:17,17 102:22 125:6 137:18 145:1 158:19 <b>govern</b> 57:11 <b>governed</b> 42:5 <b>gps</b> 25:5 <b>grab</b> 130:11 <b>grail</b> 70:16 <b>granted</b> 135:1 <b>granularity</b> 146:8 <b>graph</b> 51:16,24 52:9,12 54:15 69:7 120:9 125:9
--	--	--	--



## [graph - identifiers]

136:25 137:2,8,9 137:10,14 <b>grasp</b> 135:6 <b>great</b> 41:23 96:15 111:6 124:8 <b>group</b> 8:10 22:5 54:23 <b>grouped</b> 75:18 76:2 <b>groups</b> 49:10 50:18 <b>guarantee</b> 103:21 <b>guess</b> 8:9 10:11,16 15:9 24:4 25:23 27:1 32:6 38:3 50:1,25 55:20 56:16 57:18 61:14 68:16,21 74:11 76:16 86:25 89:7 100:4,11 103:15 104:8 105:25 106:25 107:4 119:8,8,22,24 121:2 123:4 126:5 144:3 147:3 148:12,17 <b>guide</b> 14:16,18 <b>guys</b> 10:21 12:21 18:16 23:21 43:13 44:2 46:11 48:22 63:10 66:3 94:10 99:24 105:13 109:3 130:22 131:13 153:20 154:3 156:8	<b>handle</b> 42:8 <b>handled</b> 141:7 <b>handles</b> 106:5 <b>hands</b> 86:23 <b>happen</b> 112:16 <b>happened</b> 70:10 99:1 <b>happening</b> 112:20 138:8 <b>happens</b> 52:23 <b>hard</b> 48:22 76:20 77:1 134:20 158:18 <b>hardware</b> 23:22 76:19 <b>hash</b> 24:12,13,14 24:20,23 25:1,9,10 26:12,13,15,22 84:13 <b>hashed</b> 83:9 114:14 115:23 116:8,9 <b>hashes</b> 24:21,21 <b>hashing</b> 24:15 26:10 <b>hazel</b> 45:7 <b>head</b> 60:10 84:23 97:10 <b>hear</b> 29:12 47:24 76:10 92:16 107:7 111:1 116:6 133:12 150:2 158:21 <b>heard</b> 102:1 107:22 <b>hearing</b> 1:13 2:10 6:3 156:22 158:3 159:11 <b>hearings</b> 6:5 <b>height</b> 45:7	<b>held</b> 29:16 <b>hello</b> 5:16 <b>help</b> 7:24 17:22,23 41:20 52:13 72:21 121:20 122:5 133:4 135:18 145:17 <b>helpful</b> 7:7 19:13 50:7 62:5 115:25 137:20 143:4 145:16 148:2 157:2,17 159:3 <b>helping</b> 7:9 <b>helps</b> 104:1 <b>hey</b> 120:1 <b>high</b> 69:2 70:6 101:4 120:11 140:16 <b>highest</b> 140:18 <b>highly</b> 1:9 158:4 158:10 <b>history</b> 7:1 15:21 <b>hit</b> 93:4 146:22 <b>hive</b> 33:4,13,23,24 33:25 34:4,5 35:1 36:18,21 37:25 38:4,11,19,25 39:3 39:8 53:9,10 54:11 75:4,13,25 76:18 80:24 81:2 82:14 90:19 91:4 92:6,20,24 93:2 102:15,20,24 103:2,6,7,12 110:10,15 111:10 111:14,15,16,17 112:2,5,7,20,22 113:16 114:17,24 115:7 118:25 120:4 129:8,9,13 129:16,19,23	130:6,7,10,10,11 130:12,19,21 131:5,23 132:1,3 132:24,24 143:12 143:12,20 144:14 144:19 145:3 154:18 156:5,22 <b>holy</b> 70:16 <b>hometown</b> 88:23 <b>homework</b> 14:3 <b>hope</b> 71:11 157:17 <b>hopefully</b> 6:4 10:7 63:15 100:24 157:18 <b>hoping</b> 13:7 16:11 <b>hour</b> 5:6 <b>hours</b> 145:6 <b>house</b> 25:22 141:13 <b>http</b> 97:16 <b>huh</b> 82:7 <b>human</b> 22:25 23:2 23:4,5,11,12 <b>hype</b> 136:21 <b>hypothetical</b> 59:19
<b>h</b>			
<b>hack</b> 122:10 <b>hacker</b> 4:3 <b>hand</b> 7:14 9:3,11 9:19 10:2 12:4 32:13			<b>i</b>
			<b>i.e.</b> 69:6 <b>ian</b> 4:2 <b>ianchen0</b> 4:4 <b>idea</b> 85:7 <b>identified</b> 28:13 30:3 78:20 <b>identifier</b> 94:21 105:6 111:18 120:5 <b>identifiers</b> 30:19 104:19,19 105:3 105:12 114:13 116:2,3 150:12

**[identifies - interested]**

<b>identifies</b> 155:12	95:10 146:21	<b>information</b> 12:17	<b>integer</b> 121:9
<b>identify</b> 6:7 7:24	<b>included</b> 14:8 15:5	14:1,19,22 15:4,25	<b>integers</b> 105:15
88:1 101:11	42:21 43:1,5 44:4	21:25 26:22 32:3	<b>integrate</b> 64:23
121:20	47:11 57:20,20	46:24 47:1 59:22	65:16 66:5,8,12
<b>identifying</b> 95:11	58:6 59:9,11 61:7	64:9,10 70:10	67:15 71:17 73:2
<b>ids</b> 105:14,16	79:25 80:1 85:12	75:5 78:2,3 80:23	79:6 122:15
106:5 107:12	85:24 95:4 97:3,4	83:5,7,8 85:16	<b>integrated</b> 64:21
111:15,17 121:4	99:13 122:8 147:4	86:9,11 87:16,18	86:12
130:19	<b>includes</b> 13:13	87:23 88:20,24	<b>integrates</b> 65:5,23
<b>images</b> 127:3	49:9 151:12	94:25 95:1 98:10	68:23
<b>imagine</b> 91:11	<b>including</b> 93:24	99:12 104:8,10,19	<b>integrating</b> 64:18
<b>immediate</b> 53:4	95:14,17 98:14	104:22 105:9,20	65:4
<b>immediately</b> 52:23	158:15	105:21 106:11,12	<b>integration</b> 67:12
69:19	<b>income</b> 147:24	106:21 107:3,4,20	68:12 71:18 100:1
<b>implement</b> 53:19	<b>incorporated</b>	119:21 141:21	<b>integrations</b> 71:14
124:1 125:10,19	154:20	147:16,17 158:5	71:25
<b>implementation</b>	<b>incorrect</b> 29:22,25	<b>informative</b>	<b>integrity</b> 26:3
138:12	<b>incredibly</b> 119:19	157:17	37:17 58:11 59:1
<b>implementations</b>	<b>independent</b> 27:3	<b>infrastructure</b>	93:1 96:8
138:9	27:8	52:15,19 109:12	<b>intell</b> 99:19
<b>implemented</b>	<b>index</b> 76:14	134:12 136:20	<b>intelligible</b> 98:13
79:15 124:14	135:16	138:5 146:24	98:14,19,21
<b>implementing</b>	<b>indexed</b> 75:21	<b>initial</b> 13:5 28:6	<b>intention</b> 54:8
52:17 91:22	144:12	<b>initiate</b> 94:19 95:4	<b>interact</b> 16:19,25
125:22	<b>indexes</b> 129:8	149:20	17:2 89:19 121:23
<b>implica</b> 147:6	<b>indexing</b> 136:17	<b>input</b> 41:10 84:6	<b>interacted</b> 49:13
<b>implications</b> 146:3	<b>individual</b> 42:25	<b>inputs</b> 96:2	<b>interaction</b> 50:17
147:4 149:25	43:4 44:3 59:7	<b>ins</b> 64:25	110:6
150:5	60:3 92:12 119:5	<b>inside</b> 16:25 30:20	<b>interactions</b> 16:23
<b>important</b> 18:25	125:23 141:5	117:17	49:9 50:20 86:6
133:13	149:18	<b>insights</b> 88:9	90:11 138:7
<b>importantly</b>	<b>individually</b> 115:2	<b>instance</b> 22:15,19	<b>interactive</b> 15:3
152:13,14	<b>industrywide</b>	26:5 49:12 72:19	<b>interdependencies</b>
<b>impression</b> 48:19	109:4	72:22 122:7	35:8
<b>inaccurate</b> 62:18	<b>inference</b> 151:17	137:19	<b>interdependency</b>
<b>inaudible</b> 65:11	<b>inferred</b> 13:16	<b>instruct</b> 71:16	35:9
76:9 92:14	87:1,1 93:7	<b>instructed</b> 6:15	<b>interest</b> 88:13,18
<b>include</b> 14:13	<b>info</b> 131:11	<b>instruction</b> 121:23	143:9
22:10,16 24:21	<b>inform</b> 90:24	<b>instructive</b> 157:19	<b>interested</b> 46:19
25:7,11 45:14,17	143:24 145:18	<b>instrument</b> 27:14	61:16 80:6 81:6
58:8 65:6 86:8			88:5 101:23

[interested - laser]

160:16 <b>interesting</b> 110:8 113:13 153:24 <b>interests</b> 87:17 <b>interface</b> 110:4 127:21 <b>interfaces</b> 69:6 <b>interject</b> 48:10 133:4 <b>internal</b> 38:22 82:19 105:3,6 137:8 <b>internally</b> 83:23 <b>internals</b> 76:23 <b>interpret</b> 94:25 99:3 <b>interpretable</b> 95:19 <b>interrogatories</b> 66:16 <b>interrupt</b> 60:16 <b>interrupting</b> 8:3 <b>intimately</b> 21:19 <b>intrinsic</b> 52:6 <b>intro</b> 59:19 <b>introduction</b> 59:17 67:7 <b>intrusive</b> 91:15 <b>invested</b> 65:15 <b>invite</b> 7:15,19 <b>invited</b> 6:13 <b>involved</b> 158:25 <b>involves</b> 130:13 <b>involving</b> 141:20 <b>ip</b> 49:13,14 50:15 53:7 88:22 89:18 95:3 <b>irrelevant</b> 5:9 <b>issue</b> 7:25 13:13 133:6,9	<b>issued</b> 6:6 <b>issues</b> 134:12 156:22 <b>item</b> 94:7 95:12 136:9 <b>items</b> 59:4 <b>j</b> <b>jams</b> 10:5,7 63:15 <b>job</b> 129:11,15 157:11 <b>jobs</b> 129:11 157:11 <b>jsc</b> 1:6 2:6 <b>json</b> 97:18 99:10 99:11 127:3,12,15 <b>judge</b> 13:12 88:6 88:10 158:22 <b>jump</b> 8:1 12:13 <b>jumping</b> 37:8 42:19 <b>junkyard</b> 42:3 <b>jurat</b> 160:2 <b>k</b> <b>k</b> 3:5 <b>keep</b> 14:21 49:19 56:22 60:25 64:10 65:18 93:11,16 102:5,11 109:1 143:9 <b>keeping</b> 91:15 <b>keller</b> 3:9 <b>kellerrohrback.c...</b> 3:12,13 <b>key</b> 76:6 105:17 118:10 144:15 <b>kick</b> 11:24 13:23 <b>kicks</b> 113:2 <b>kidding</b> 100:19 <b>kids</b> 57:24	<b>kind</b> 29:13 30:6,18 38:17 59:14 76:19 77:20 97:19 108:9 111:18 121:20 <b>kindly</b> 5:9 <b>kinds</b> 57:7 79:17 108:12,13 <b>knew</b> 20:5 <b>know</b> 8:10,21,23 10:17,21,24,24,25 11:2,4,4,7,11,22 14:1,11,20,22 17:12 19:16,19,19 20:3 21:17 22:4 22:12,16,21 23:17 23:21 24:6 25:4 25:13 31:17 34:17 35:22 36:2,6 38:21 40:9 43:6 43:14 44:1,9,24 45:1,3,5 48:5,12 48:16 51:23 52:6 52:8 54:24 55:6,8 56:9 57:8 58:2,7 58:11,15,24,25 59:2 60:14 61:8 61:20 62:11 64:14 68:17 71:3 73:14 75:11 76:22,24 77:21,23,24 78:2 78:15 80:10,21 82:16,18 84:10,17 84:19 85:5,6 87:22 88:17,19,23 89:3 90:14,19 93:23 95:1 97:25 99:3,13,24 101:7 101:21 102:6,21 103:3,3,25 104:8,9 105:2,13 106:8,20 108:2,7,9 113:5,6	115:16,25 116:12 116:12 118:6,19 118:21 121:4,22 127:14 128:7,9,10 132:19 133:14,21 134:4 135:4 137:3 139:20,23,23,25 140:1 141:14,16 141:19 142:13,15 143:4 144:17 147:14 148:19,20 149:14 151:2,5 152:11,23,25 153:21 157:10 <b>knowledge</b> 42:23 85:17 86:1 119:25 <b>knows</b> 8:20 52:8 96:17 107:22 140:2 <b>ko</b> 4:8 <b>kutscher</b> 3:16 <b>l</b> <b>lack</b> 37:25 <b>language</b> 34:7,13 87:20,21,21 89:2,5 122:11,12 126:3 <b>languages</b> 39:16 87:23 <b>large</b> 17:16 67:25 119:20 127:5 129:17 130:3 <b>largely</b> 6:12 <b>larger</b> 76:3 127:10 <b>laser</b> 76:4,14,21 76:23 103:2,4,12 129:8,8,15,15,18 129:23,25 130:14 130:25,25 135:15 135:19 136:3,11 136:17,19,25 137:2,3,10,12,14
---	--	---	--



[laser - magnitude]

144:5,11 <b>laser's</b> 136:23 138:2 <b>latency</b> 129:1 <b>latest</b> 127:7 <b>laufenberg</b> 3:10 <b>lawyer</b> 23:7 <b>lawyers</b> 23:6,9 113:12 157:11 <b>lay</b> 8:2 <b>layer</b> 106:9 108:8 125:5,8,20 126:23 <b>layers</b> 144:23 <b>lead</b> 10:23 <b>learning</b> 40:14 <b>leave</b> 48:19 67:12 97:8 <b>legal</b> 4:2 150:23,25 151:1,18 <b>lengthy</b> 77:13 <b>lens</b> 83:4 <b>lesley</b> 3:4 <b>letters</b> 13:18 <b>letting</b> 153:20 <b>level</b> 44:21,22 45:23 49:12 67:21 68:16,25 69:2 70:6 75:19 101:4 120:11 140:16,18 146:8 <b>license</b> 65:9 <b>life</b> 74:17,21 146:17 <b>lifetime</b> 114:6 <b>light</b> 68:9,10 <b>liked</b> 52:24 53:2,5 <b>likelihood</b> 153:14 <b>liking</b> 88:15 <b>limit</b> 54:25 127:7,7 <b>limited</b> 61:1 72:21 88:7 107:4	<b>line</b> 35:18 68:18 70:2,4 74:10,20,23 74:24 77:8 78:6,8 78:11,12,13,14,18 78:20,24 79:7,7,10 79:13 80:7,12 81:6,7 82:2,11,21 86:8,11 91:2 114:14 <b>lines</b> 10:21 <b>link</b> 43:18 <b>linked</b> 16:15 28:10 <b>list</b> 17:16 35:15 39:17,19,20 41:16 53:25 65:2,3 67:16,22 73:21 74:1,4,6,8 86:1 87:17,25 91:6,6 114:17 119:13,16 127:23 <b>listed</b> 68:17 <b>lists</b> 49:19 82:24 <b>lit</b> 54:1 <b>literally</b> 41:24 66:17 <b>litigation</b> 1:7 2:7 6:11,25 135:5 <b>little</b> 18:20 72:23 112:4 144:3 159:4 <b>live</b> 15:3,25 31:22 89:17,21,23 <b>lives</b> 121:3 <b>llp</b> 3:9 <b>load</b> 22:7 136:13 136:21 <b>locale</b> 88:25 <b>locality</b> 128:14,14 128:15 <b>locally</b> 128:21 <b>location</b> 16:8 30:9 49:18 84:18 87:9	87:10,11,12,15 88:11,22 89:16,21 104:12,13,14 105:10 <b>locked</b> 57:23 <b>loeser</b> 3:10 <b>log</b> 16:7,9,10,12,16 16:16,20,21,25 17:2,6 18:2,21,22 19:15 20:9,13 21:6 22:11 27:5 46:22 51:20,21 53:9 97:18 155:12 155:13,16,17,17 <b>logged</b> 32:19 50:16,21 155:9,11 <b>logging</b> 16:23 <b>logic</b> 125:11 <b>login</b> 47:23 <b>logs</b> 33:22 <b>long</b> 11:13 15:20 17:17 26:16,17 30:12 41:16 56:17 56:22 110:9 111:9 112:11 120:4 129:1 <b>longer</b> 65:13 66:5 66:7,9,14 67:18 <b>longest</b> 133:16 <b>look</b> 15:25 41:23 44:5 62:1 64:13 76:7 94:15,16 101:5 104:14 115:3,5 117:17 123:5 128:3 130:14 131:5 133:3 139:17 141:20 159:3 <b>looked</b> 35:15 71:20 95:8 148:14 149:9 151:21	<b>looking</b> 12:18 13:3 14:2 15:17 19:2 54:6 55:15,16,18 61:24 75:12 92:8 92:9 94:12 102:17 110:5 124:4 130:9 131:8 148:24 <b>looks</b> 32:4 54:1 <b>lookup</b> 76:25 130:12 <b>lookups</b> 129:24 <b>loop</b> 23:7,7 133:22 133:24 134:3 <b>loose</b> 19:5 <b>lose</b> 5:6,11 108:21 <b>losing</b> 56:7 <b>lost</b> 76:12 99:17 <b>lot</b> 7:1,2 14:5 19:21,24 32:2,4 50:2 61:25 67:8 98:10 113:11 117:16 126:24 132:14 134:11 146:16 154:16 156:15,16 <b>lots</b> 42:13 66:17 <b>loud</b> 40:23 <b>low</b> 15:23 148:16 148:16 <b>luckily</b> 14:5 <b>luluemon</b> 94:17 95:24 96:15,20 149:6 <b>lweaver</b> 3:7
<b>m</b>			
<b>mach</b> 24:13 <b>machine</b> 24:11,20 160:12 <b>machines</b> 136:20 <b>magnitude</b> 135:7			

**[majority - mind]**

<b>majority</b> 120:20	66:10 67:17 71:3	43:12,16,20 47:19	<b>meeting</b> 30:3
<b>making</b> 7:8,8	91:13,14,19,22,24	48:1,11 50:3,4	<b>meg</b> 25:13,14
21:11 117:1 156:2	<b>martie</b> 3:16	51:5 55:1 56:8,24	<b>memcache</b> 129:6
157:16	<b>master</b> 1:11 6:2	61:21 65:7 66:16	135:19,22,25
<b>mana</b> 11:20	10:10 11:10 19:9	72:7,9 73:14 78:7	136:13,20 137:4,4
<b>manage</b> 108:8	38:13 60:12 62:10	79:15 89:9 91:5	144:11
123:22	71:12 73:20 76:11	94:2 100:2,4,5,9	<b>memory</b> 76:24,25
<b>managed</b> 41:13	77:19 93:10 96:22	100:10 104:14	<b>menlo</b> 4:4
<b>management</b>	98:17 101:6	106:12,13,23	<b>mental</b> 104:1
42:15 135:13	105:24 106:22	107:12 119:18	<b>mentioned</b> 41:20
<b>manager</b> 140:12	107:1,7 123:18	128:25 132:13	<b>messages</b> 61:7
140:13	130:16 132:21	133:13,14 134:24	<b>messenger</b> 62:4
<b>manifold</b> 120:24	133:3 145:2	135:3,8 139:9,17	<b>metadata</b> 25:3
121:3,6 124:22	147:19 148:18	143:16,18,19	<b>method</b> 119:9
127:1,2,10,13,15	158:14	144:1 146:7	<b>methodically</b>
127:21,24	<b>masterful</b> 56:8	147:21 148:1,6,12	48:23
<b>manifolds</b> 120:23	<b>match</b> 52:19 83:8	149:7 150:25	<b>michael</b> 4:7
<b>mann</b> 4:7	104:24 113:25	151:15 152:9,10	<b>michelle</b> 1:18 2:10
<b>manner</b> 77:11	114:4 126:19	152:10 155:4	5:7 160:24
<b>manual</b> 133:19	<b>matched</b> 75:20	156:13	<b>microsoft</b> 138:24
<b>manually</b> 132:17	85:19	<b>meaning</b> 99:25	139:24 141:2,10
<b>map</b> 81:11 94:24	<b>matches</b> 87:19	112:1 128:14	141:21 142:18
95:8 131:19	<b>matching</b> 77:13	<b>meaningful</b> 23:18	<b>mike</b> 12:6 57:23
132:17 149:20,20	84:14 116:4	35:10 59:12	84:11 115:4
154:9	<b>material</b> 7:20	149:22	120:15 130:22
<b>mapped</b> 21:2	<b>matter</b> 134:19	<b>meaningless</b> 59:24	133:14 144:17
108:14 110:10	138:16 151:12,20	<b>meanings</b> 151:18	145:1
111:10 112:2,13	<b>mayur</b> 9:13 17:22	<b>means</b> 30:5 53:3	<b>mike's</b> 91:10
113:15 133:21	33:8 67:2 74:20	63:17 71:2,3 79:7	<b>mike.com</b> 84:11
149:19,22	84:24 90:15	99:18 151:17	<b>milan</b> 1:18 2:10
<b>mapping</b> 44:10,12	100:12 110:15	<b>meant</b> 15:2 16:16	160:24
44:12 106:5	120:15 121:1	16:22 17:1 67:10	<b>mill</b> 3:17
108:14 109:9,17	122:21 128:2	111:5 127:4	<b>million</b> 39:18
110:9 111:9,23	130:22 137:18	<b>measured</b> 79:14	131:13,14,15
112:1,11 113:1	152:23	<b>measurement</b>	<b>millions</b> 10:21
116:3 131:1,2,7	<b>md</b> 1:6 2:6	72:21 116:15	67:25 115:3,11,18
<b>mappings</b> 112:24	<b>mdl</b> 1:5 2:5	<b>measurements</b>	131:10,10
<b>march</b> 1:15 2:13	<b>mean</b> 20:5 23:21	72:20	<b>millisecond</b> 18:19
5:1 160:19	24:10 31:1 32:24	<b>mechanism</b> 89:4	<b>milliseconds</b> 18:18
<b>marketplace</b>	35:18 38:16 39:18	<b>meet</b> 59:4 140:14	<b>mind</b> 61:17 144:10
59:16,16,18,20,21	40:3 41:1 42:22		

[mine - ofa]

<b>mine</b> 15:21 19:22 19:25 89:20 148:14 151:5 <b>minute</b> 63:8 117:7 <b>minutes</b> 5:6 49:23 63:11,11 <b>missed</b> 158:1 <b>missing</b> 132:16 <b>mission</b> 3:21 <b>misunderstanding</b> 108:2 <b>mittchell</b> 5:22,23 9:18,21 10:4,5,7 31:24 43:7 45:12 45:19 46:1,10,15 62:19 63:1,3 94:1 116:22 142:8,12 143:2 145:23 151:10 <b>mittchell's</b> 45:4 <b>mix</b> 89:13 <b>mkutscherclark</b> 3:18 <b>mobile</b> 23:21 24:7 56:17 <b>model</b> 23:15,19 104:1 138:7 <b>models</b> 23:15 145:18 <b>moderator</b> 10:7 63:16 <b>modified</b> 97:16 102:6 <b>moment</b> 149:12 <b>money</b> 81:18 <b>month</b> 7:3 <b>morning</b> 5:20,23 5:25 <b>move</b> 7:10 28:17 37:6,10 93:10 134:5,6,12,13	<b>moved</b> 36:4 76:4 <b>moving</b> 82:23 110:7 113:8 <b>multiple</b> 23:9,14 23:15 56:14 101:22 <b>music</b> 90:6 <b>mute</b> 116:21 121:15,17 <b>mysql</b> 54:15 126:8 126:10,17,20,20 126:25 127:8,8,11 127:21 137:6  <b>n</b>  <b>nailed</b> 51:13 <b>name</b> 27:21 40:3 97:9,20,24 98:4 104:10,10,22,23 105:5,8 106:13,13 118:11,12,13 120:3,5,6 160:19 <b>named</b> 6:5,10 <b>names</b> 94:20 97:6 97:6,7 118:16 <b>narrow</b> 88:9 <b>nature</b> 6:15 <b>near</b> 145:5 <b>nearly</b> 119:8 153:25 <b>necessarily</b> 69:16 70:25 108:15 149:2,7 <b>necessary</b> 6:5 <b>need</b> 5:7 14:5 17:17 18:9 20:12 41:1 53:1 62:1 63:10 75:21 83:25 84:23 90:23 92:22 92:25 93:2,9 123:19 126:13 140:15 144:22	146:13 150:6 152:1 153:3 154:15,22 156:11 157:21 <b>needed</b> 27:14 28:2 28:16 38:6 108:16 109:14 <b>needing</b> 61:20 <b>needs</b> 30:22 80:17 128:19 <b>neither</b> 160:15 <b>net</b> 66:15 <b>netflix</b> 66:4,11 138:24 139:24 141:1 142:18 <b>never</b> 22:20,23 23:6 46:23 117:22 144:10 <b>new</b> 17:18 36:15 41:18 65:17 66:15 91:10 143:6 155:14 <b>news</b> 41:9 49:15 <b>nice</b> 113:13 <b>night</b> 131:14 <b>nobody's</b> 99:6 <b>node</b> 121:13,18,21 122:3,9,11,16,19 123:8 124:19 126:1,2 <b>non</b> 24:13,14,22 46:13 110:19 <b>normal</b> 24:5 <b>north</b> 24:9 <b>northern</b> 1:2 2:2 <b>noted</b> 159:6 <b>notice</b> 18:17 20:9 27:14 36:9 <b>noticed</b> 148:24 <b>number</b> 38:17 52:14 55:17 67:1	67:4,9,24,25 69:7 72:21 75:7 100:22 104:11,23 105:6 125:16  <b>o</b>  <b>oakland</b> 3:6 89:23 <b>oath</b> 160:10 <b>object</b> 6:14 43:16 43:17 45:6 50:24 136:7 <b>objectionable</b> 58:13 <b>objectives</b> 72:8 <b>objects</b> 117:14,25 119:5 126:9 127:3 127:24 128:8 <b>obligations</b> 29:6 30:3 74:5 <b>observation</b> 54:7 58:15 <b>observe</b> 50:2 55:4 <b>observed</b> 49:3,4 49:21 50:23,24 51:2,5 52:2 53:7 53:16,20,21,25 54:8,9,18,21 55:3 55:5,24 56:14,18 57:5,7,19 58:5,10 58:19 59:9,10 60:4 64:6 89:18 89:18 <b>obtained</b> 13:14 <b>obviously</b> 45:21 59:13 124:23 <b>occasionally</b> 145:17 <b>occurred</b> 112:24 <b>occurs</b> 84:14 <b>ofa</b> 38:16,19,20 39:21,22 40:22 41:22 48:13 66:19
---	--	--	---

[ofa - party]

74:19 75:2,9 77:12,23 78:1,2,3 79:21,25,25 80:10 81:3,7 82:21 137:19 <b>offer</b> 65:14 66:9 66:14 70:23 109:3 130:6 138:11 <b>offered</b> 70:24,25 71:1,2,23 <b>offers</b> 137:11 <b>offhand</b> 118:21 <b>officially</b> 6:22 7:11 <b>oh</b> 8:17,18 11:17 18:25 20:1 26:19 37:6 40:19 60:21 73:9 121:2 132:18 155:25 <b>okay</b> 5:19 8:13,18 10:20 12:3,11 14:20 15:15 16:6 17:4 18:12,21 20:4 26:19 27:1 30:8 31:3,9,14 35:21 40:7,22 42:25 44:3,17,25 46:10,15,17 48:4,9 48:18 49:3 52:2 57:14 58:4 60:21 62:8 63:8 70:1,18 77:15 78:17 79:1 79:23 86:24 92:17 93:17 94:15 96:5 96:6,16 101:16,18 101:20 104:3 105:11 107:18 110:14 111:1 112:10,14 113:8 116:9,11 117:3 120:24 126:1,7 128:6,24 131:25	132:19 135:6 136:23 138:15,16 144:12 145:14,25 149:23 152:1 158:12 <b>old</b> 141:12 143:7 <b>older</b> 34:21 <b>once</b> 38:4 90:17 91:8 112:24 127:6 136:5,12 <b>ones</b> 33:15 35:15 46:9 51:17,18 59:12 65:13 67:18 126:4 <b>online</b> 115:7 <b>open</b> 59:16,21 63:18 <b>opening</b> 150:16,19 151:7 <b>operate</b> 101:22 108:23 121:22 <b>operated</b> 103:6 <b>operates</b> 128:11 132:12 <b>operating</b> 112:25 121:14 <b>operation</b> 130:7 <b>operationally</b> 41:23 <b>opportunity</b> 93:12 <b>opposed</b> 16:18 <b>optimized</b> 144:20 <b>option</b> 30:5 46:6 148:3,16 152:20 <b>options</b> 147:24 151:22 152:10,19 <b>order</b> 6:6,9 18:8 28:9,11 53:1 70:13 121:22 144:23 155:16	<b>original</b> 24:25 25:11,12 28:14,15 125:14 <b>orm</b> 125:5,8,20 126:23 <b>ouage</b> 12:1 <b>outcome</b> 52:17,20 <b>outcomes</b> 138:13 <b>outline</b> 14:3 30:18 49:7 115:22 <b>outlined</b> 30:2 89:25 <b>output</b> 35:24 36:6 37:22 41:11 44:13 44:14 77:17 133:23 <b>outputs</b> 153:8 <b>outside</b> 16:18 122:17 132:1 <b>overall</b> 62:10 141:4 142:7 <b>owner</b> 81:13 <b>owns</b> 108:18 109:8 109:12	<b>pain</b> 61:15 132:5 <b>palm</b> 56:24 <b>palo</b> 3:17 <b>paperwork</b> 42:14 <b>paragraph</b> 93:23 150:16,20 151:7 <b>params</b> 97:19 <b>pardon</b> 37:8 <b>part</b> 15:22 20:18 20:24 21:13,24 28:7 29:1,3,4 30:11 42:6 44:10 50:14,16,17 57:2,3 71:17 72:18,24 77:14 79:22 81:5 81:7 83:17,21 85:14 90:24 108:7 108:7,22 116:6 122:16 133:13 142:7 143:17,21 153:11 154:21 <b>participate</b> 157:12 <b>particular</b> 11:6 17:12 43:9,10 68:19 82:11 88:4 140:8 <b>particularly</b> 33:21 33:22 <b>parties</b> 13:15 39:11 64:11,15 67:20 113:18 114:4,10,12 115:24 156:7 160:17 <b>parts</b> 36:22 <b>party</b> 38:23,24 39:2 47:12,21,21 47:23 64:16,17,18 68:16,24 71:16 72:1 74:5 86:12 86:15 88:11,14
		<b>p</b>	
		<b>p.m.</b> 5:2 117:9,10 159:11 <b>pacific</b> 2:13 <b>page</b> 3:17 14:16,17 14:18 48:5 49:16 52:22,24,25 53:2,3 54:22 79:4,12 83:16 88:16 94:18 94:19 111:6 121:11 145:16,20 145:21,23 149:20 150:17 151:8,10 <b>pages</b> 21:25 49:10 50:18 51:23 52:7 55:6,7 88:16	

[party - please]

89:24 99:7 136:24 137:13 138:3,6,7 140:20 141:9 146:21 149:17 155:3,19 <b>passed</b> 113:19,21 113:22 156:18 <b>password</b> 110:2 <b>patching</b> 65:20 <b>patel</b> 5:15,16 8:15 9:10,13 17:22,25 18:6,14,24 19:14 24:15,17,23 25:8 25:10,19 26:9,19 26:21 31:19 33:8 33:10,19 34:1,15 34:25 35:5,20 36:1,14,20,24 37:2 37:4,6 38:8,10,15 39:20,24 40:19,21 41:22 46:4 48:12 51:10,11,12 52:6 59:11 60:8 62:19 62:24 64:20 65:7 65:10,12,22 66:18 66:25 67:3 70:4 75:1,17 76:2,11,17 77:6,10,16,20,25 78:22 80:15 81:1 81:21 82:7,18 84:25 85:5 86:3 88:4,13 95:19 96:6 97:5,12 98:8 98:16 104:13 106:12,16 110:12 110:18 113:17,21 113:25 114:3,11 114:19,23 116:15 118:21 121:2,15 122:22 123:24 124:20 125:6	126:2 128:3,10 132:8,11 133:24 134:3,10 137:22 145:20 146:7,12 147:6,21 148:1,6 148:18 152:25 154:7,25 155:7,25 156:10 <b>path</b> 28:9 132:23 <b>pattern</b> 69:5 <b>patterns</b> 146:23 <b>pdf</b> 145:23 <b>people</b> 19:12 28:23 35:6 49:17 56:24 63:5,13 65:14,17 66:15 73:2 113:11 116:12 123:25 148:5 <b>percent</b> 11:17 19:3 92:8 110:5 144:16 158:9 <b>percentage</b> 148:17 <b>perfect</b> 26:20 31:25 93:21 <b>perform</b> 88:15 89:11 <b>performs</b> 127:8 <b>period</b> 62:10 135:25 143:5 <b>periods</b> 57:2 <b>permission</b> 71:4 <b>perpetuity</b> 67:13 <b>person</b> 53:8 74:11 91:18,21,24,25 92:12 95:6 96:17 99:19 100:1 104:9 124:8,18 <b>personal</b> 94:16 <b>personally</b> 73:13	<b>persons</b> 67:10 <b>perspective</b> 23:17 26:4 38:3 44:16 50:23 58:11 59:1 79:19 89:14 144:22 <b>phone</b> 12:2 56:10 88:11 104:11,23 105:6 115:4 131:8 131:9,10 155:15 <b>phones</b> 56:8 <b>photo</b> 22:21,22 24:24 25:4,9,10,11 25:12,13,14,17 26:11,13,15 121:5 121:11 123:2,3,7 124:7,10 125:3 <b>photos</b> 15:23 25:4 26:18,24 123:25 <b>phrase</b> 141:13 <b>phrased</b> 111:2 <b>pick</b> 89:2 146:23 <b>picture</b> 22:20 23:22 24:23 39:15 <b>piece</b> 18:23 91:20 92:10 109:12 122:24,25 124:4,9 126:15 146:24 154:22 <b>pieces</b> 48:23 52:15 52:19 134:16 138:5 <b>pii</b> 70:11 114:14 114:14 115:23 116:8 <b>pipe</b> 33:6 36:11 <b>piped</b> 75:25 76:1 <b>pipeline</b> 33:7,10 33:14 35:22 37:23 38:19 39:21,22 40:3 41:7,17 52:7	69:14 <b>pipeline's</b> 35:24 <b>pipelines</b> 32:21 33:6,17,19,24,25 34:5,13 35:3 36:4 37:12 38:15 39:18 39:21 40:18 52:4 82:13 <b>pipes</b> 75:15 <b>pixel</b> 64:24 97:14 <b>place</b> 39:23 50:13 89:16 94:11 95:21 105:19 141:22 142:24 143:5 160:8 <b>placed</b> 128:23 160:10 <b>places</b> 55:17 140:12 <b>plaintiff</b> 6:5,10 93:6 126:6 153:15 156:24 158:15 <b>plaintiffs</b> 3:3 7:19 7:23 32:13 86:21 <b>platform</b> 13:14,15 13:17 34:21 39:9 40:22 42:20 46:22 47:12 49:1,2,5 50:9 51:7 60:25 61:1,5 64:5,7,12 64:15 67:5,6 68:4 68:25 70:8 72:1 73:13,24 80:14 86:15,20 89:13 90:1,2,9,13,21 109:6 113:14 120:19,21 140:21 152:19 155:2 <b>please</b> 11:1 24:17 135:10
---	---	--	--

**[plug - product]**

<b>plug</b> 64:25	<b>power</b> 12:1	119:9 140:19	58:16 67:8 68:5
<b>plus</b> 15:11,13	<b>practical</b> 56:6	143:23 155:10	73:19 77:13 79:11
136:20	68:15	<b>principles</b> 58:8	81:17 83:18 92:11
<b>point</b> 10:19 11:20	<b>precise</b> 110:16	147:1	100:15 109:8,12
11:23 19:17 22:5	112:5	<b>printed</b> 118:5	113:3,4 133:13
27:4 30:21 39:4	<b>precisely</b> 93:5	<b>prior</b> 27:20 64:10	135:13 137:19
45:4,4 47:23	<b>precomputation</b>	132:2 133:11	138:18,19 139:25
50:20 61:16 73:17	130:3	148:25 160:10	140:2,7,9,10,17
76:24 84:20 90:17	<b>preexisting</b> 131:6	<b>privacy</b> 1:6 2:6	141:9,16,18,23,25
93:14 98:5 101:22	<b>preface</b> 19:1	20:19,20,22,24,25	142:9,13,23,24
117:2 127:10	<b>preference</b> 154:12	21:11,23 27:6,13	143:5 145:7
129:8 134:18	<b>preferences</b> 148:7	27:15,19,21,22	<b>processed</b> 50:21
145:1,5	<b>preferred</b> 87:21	28:3,4,9,21 29:4	69:13,23 74:14
<b>pointer</b> 30:14	89:16	30:3 42:5,10,14	82:12,13 84:15
<b>pointers</b> 30:22	<b>prepare</b> 157:14	44:12 54:5,16	<b>processes</b> 75:15
<b>pointing</b> 30:17	<b>prepared</b> 10:13	56:2 57:2 72:24	77:3 83:11 112:19
<b>points</b> 40:5 91:7	11:14,15	73:4 108:4 122:16	<b>processing</b> 33:11
142:17 146:5	<b>prescriptive</b> 20:10	122:23 123:12	33:11,14 34:5
<b>policies</b> 57:11	<b>present</b> 4:6 62:11	124:10,23 125:2,4	<b>produc</b> 17:15
<b>polls</b> 31:8	74:21 114:25	125:11,22 140:24	<b>produce</b> 18:8
<b>poorly</b> 127:9	115:1	142:1,15	26:13 39:25 40:17
<b>pop</b> 91:12 92:1	<b>presented</b> 12:15	<b>private</b> 73:10	40:20 41:8 51:20
<b>popped</b> 91:18	<b>preserving</b> 125:5	137:7 142:21	51:21 89:3 149:2
<b>populate</b> 30:22,23	<b>pressure</b> 100:19	<b>privileged</b> 6:15,18	<b>produced</b> 6:10
129:15,18	<b>presto</b> 33:15 34:6	<b>pro</b> 140:4	77:12
<b>populated</b> 15:6	34:16,23,24,25	<b>proactive</b> 136:11	<b>producing</b> 51:21
<b>portal</b> 90:10	41:10	<b>probably</b> 72:2	77:18
<b>portion</b> 120:1	<b>pretend</b> 56:10	81:23 89:1 100:6	<b>product</b> 16:16
<b>position</b> 48:8	58:1	132:22 146:13	20:21 21:5,6 28:6
<b>possibility</b> 102:18	<b>pretty</b> 53:22 54:2	153:17	28:7,10 30:7 39:9
<b>possible</b> 19:14	56:7 102:1 118:3	<b>problem</b> 124:22	49:17 50:14,25
29:19,20 40:10	<b>previous</b> 132:14	<b>problems</b> 6:24	51:2,8,13 52:5,5
45:15,22 76:6	145:7	<b>proceedings</b>	52:16,17 53:14,16
114:16,19 115:16	<b>previously</b> 110:17	157:12 160:7,9,11	53:19 54:10 55:9
128:4 138:1	155:9,24	<b>process</b> 14:10 20:8	55:23,23 56:2,12
<b>post</b> 125:12,14	<b>primarily</b> 34:25	20:10,25 21:4,21	56:17,17,23 57:9
<b>posts</b> 61:6 119:6	36:14,20 52:12	21:23,24 26:6	57:10,11 58:21,23
<b>potentially</b> 34:15	120:20 155:8,11	27:16,22 28:1,3,24	59:13,15,25 60:4,6
39:25 88:23	<b>primary</b> 87:9,10	29:1 30:17 35:10	61:22,22 62:2,2,4
129:17	87:12,15 94:23	37:16 42:14 44:6	69:9,11,17,20 71:7
	105:19 110:4	45:3 56:20 58:12	71:8,13 74:18,18



**[product - questions]**

77:5,6,17 79:21 81:1 83:24,25 84:6,7,16 86:12,13 90:12,24 91:9,16 100:16,17,20 108:17 118:25 119:1 120:2,11 125:17 135:12 138:9,13,18,19,24 140:9,12 141:24 142:4 143:13,17 143:18,24 144:21 144:24 145:11 146:19 148:20 149:17 <b>production</b> 17:8 17:11,14,15,18 18:23 19:16 32:22 37:19 69:18 76:5 102:25 103:1 134:23 145:9 <b>products</b> 15:1 21:3,3 51:3 55:12 55:12 57:11 61:23 65:3 68:3 69:8,18 69:19,20,22 71:11 71:15,25 72:12 73:13 79:5 136:24 137:13 138:3,6,11 138:22,23,25 139:2,5,14,16 140:4,6,7,16,20,21 141:5 142:14 144:15 <b>professionals</b> 21:10 <b>profile</b> 1:6 2:6 49:10,17 88:24 89:23 <b>profiles</b> 50:19	<b>program</b> 20:19,24 29:8 30:4 42:5 57:2 72:25 140:11 142:7 <b>programming</b> 122:10,11,17 <b>programs</b> 57:5,8 72:15 <b>project</b> 11:20 140:11,12,13 143:22,23 147:2 <b>projects</b> 139:5 <b>promise</b> 100:23 <b>promote</b> 71:2 <b>prone</b> 125:18 <b>proper</b> 73:3 <b>properly</b> 67:11 <b>protected</b> 107:3 <b>provide</b> 7:20 24:25 26:11 40:17 46:4,6 49:20 59:21,22 64:11,15 66:12 70:13 79:8 79:20 82:24 97:22 98:1 153:7,13,20 154:13 <b>provided</b> 7:7 21:2 21:24 42:20,21,22 42:24 44:13 45:25 46:3,5,12,13,20 47:3,6,7,8,9,12,16 47:18 48:20,25 64:9 67:24 89:15 96:14 125:10 139:24 <b>provides</b> 96:8 <b>providing</b> 46:2 68:24 99:19 152:15 <b>pseudonymization</b> 109:5	<b>pst</b> 159:11 <b>public</b> 71:19 97:13 97:13 99:9 124:12 <b>pull</b> 128:4 133:18 134:21 <b>pulled</b> 75:1 88:25 121:6 133:9 <b>purchase</b> 94:19 97:8 <b>purchased</b> 46:24 96:1 <b>purchases</b> 70:10 <b>purpose</b> 6:3 18:15 52:16 135:20 155:10 <b>purposes</b> 101:25 116:4 155:8 <b>pursuant</b> 6:9 <b>put</b> 19:7 20:15 25:17 26:4 41:14 43:17 45:8 46:22 68:9,10 81:10 88:19,23 89:17 90:5 101:9,10 129:23 133:23 149:6 156:2 <b>puts</b> 95:6 <b>putting</b> 16:4 106:14 136:24 153:19 <b>python</b> 35:6 <b>q</b> <b>quality</b> 96:9 <b>quantity</b> 55:15 <b>queried</b> 103:8 <b>queries</b> 129:14 145:8 <b>query</b> 34:12,20 35:12 86:24 97:19 103:1 125:20 129:9,16,23 133:2	136:21 144:4 <b>querying</b> 144:8 <b>question</b> 7:14 8:24 11:7 12:21 13:24 13:24 17:23 18:1 18:1 26:10 27:1 32:1,6,7,11,14,15 32:16 38:17 40:12 43:3,19,22 46:18 48:11,24 49:23 51:19 55:2,4 56:4 57:18 58:5 60:18 60:20 61:17,18 62:9 64:8 66:18 66:20,22 67:14 68:15,19 71:6,10 71:12 74:16 75:11 80:11 82:10,14 84:3 93:23 95:13 96:10 100:4,9,11 100:22 101:9,23 102:14 103:13,16 104:7,18 105:2,9 105:24 111:7 112:8 114:20 115:20 116:24 117:12,13 119:3 137:23,23 138:17 138:17 139:12,13 141:6,7,8,17 143:12 144:7,11 146:2 147:4,16 149:12 150:14,22 150:23 151:4 152:22 155:20 156:1 <b>questions</b> 6:13,16 6:17 7:3,4,13,15 7:16,18,20,24 8:14 11:14 12:15,16,19 32:12,17 47:5
--	---	---	--

## [questions - reporter]

60:10,11 61:1,11 61:12 86:21 93:7 93:13 101:10,17 103:4,11 126:6 135:12 143:7 153:3 154:12,14 156:25 157:15 158:17 <b>quick</b> 100:18 <b>quicker</b> 76:24 <b>quickly</b> 11:24 130:11 137:5 <b>quite</b> 14:3 20:10 109:13 118:7 121:13 123:23 126:10 129:17 134:15 <b>quotes</b> 102:16	24:11,14,20 <b>reading</b> 56:5 139:6 <b>reads</b> 144:21 <b>real</b> 15:17,19 18:9 26:7 53:4 76:7 90:24 100:4,9 119:10 123:20 145:5 154:2 <b>realize</b> 6:25 7:1 19:5 116:24 <b>really</b> 13:24 17:17 25:23 34:8 35:12 53:11 96:3 107:8 109:14 110:16 118:2 121:4 126:23 130:6 133:4 138:15 139:8 144:25 157:13 <b>reason</b> 46:20 85:10 155:18 <b>reasonable</b> 115:17 <b>reasons</b> 22:9 61:14 93:24 95:13,17 98:12 <b>rebecca</b> 88:17 <b>receive</b> 83:8 84:13 113:17 114:9,11 <b>received</b> 94:21 95:4 115:24 <b>receives</b> 80:16 84:12 <b>receiving</b> 62:12 <b>recertify</b> 20:21 42:13 <b>recess</b> 63:25 117:8 <b>recognize</b> 6:22 7:2 7:11 13:2 <b>recommend</b> 37:19 156:9	<b>recommendable</b> 37:17 <b>recommendation</b> 21:10 <b>recommendations</b> 90:6 <b>record</b> 5:4,13,14 6:22 7:11 13:1,5,9 23:25 29:15,16,17 63:20,21 64:1,2,3 64:4 72:23 117:7 117:9,10,11 157:23 159:6,8 160:11 <b>recorded</b> 52:25 75:8 <b>recycled</b> 114:7 <b>red</b> 141:14 <b>redo</b> 38:5 <b>reduce</b> 102:9 <b>refer</b> 30:23 44:2 45:5 47:10 90:8 105:17 <b>reference</b> 16:7 147:14 <b>referenced</b> 30:9 80:7 <b>referred</b> 30:16 140:23 <b>referring</b> 13:21 17:12 30:20 41:22 77:25 86:9 147:23 151:13 155:1 <b>regard</b> 115:23 <b>regarding</b> 6:5 12:18 13:15 42:14 <b>register</b> 68:10 <b>regret</b> 35:22 <b>regular</b> 20:21 28:8 157:13	<b>reiterate</b> 149:11 <b>reject</b> 30:5 <b>rejoins</b> 10:4,6 <b>related</b> 120:2 <b>relates</b> 7:24 <b>relating</b> 6:9 <b>relationship</b> 52:25 <b>relative</b> 43:16 160:16 <b>relativity</b> 135:6 <b>released</b> 30:7 <b>relevant</b> 7:21 69:11 <b>relief</b> 135:1 <b>reload</b> 53:3 <b>remediation</b> 30:6 <b>remediations</b> 28:14 <b>remember</b> 34:10 94:8 127:7 <b>remind</b> 5:9 <b>reminding</b> 8:6 <b>remote</b> 1:13 2:10 57:25 <b>removed</b> 102:9 <b>repeat</b> 43:3 74:8 92:7 107:25 111:7 135:16 151:3 <b>repeated</b> 35:16 137:23 <b>repeatedly</b> 125:19 <b>replacement</b> 110:7 111:17 <b>replicated</b> 128:19 <b>reported</b> 1:18 121:17 <b>reporter</b> 2:11 5:17 5:18 8:7,9,11,13 8:18 9:2,7,10,15 9:18,23 10:1 12:3 12:8 29:11 63:20
<b>r</b>			
<b>raise</b> 7:14 9:3,11 9:19 10:2 12:3 32:13 <b>ran</b> 82:16 131:15 <b>random</b> 95:20 <b>randomized</b> 136:14 <b>range</b> 154:19 <b>ranges</b> 114:12 <b>raw</b> 32:19 33:22 44:19 125:9 <b>reactive</b> 135:22 <b>read</b> 12:16 13:8 14:5,6 17:1 20:7 23:6 27:2 32:3,16 48:1,2,6,7 63:5 66:2 73:2 94:3 129:5 132:13,20 133:6 145:10 146:23 157:3 <b>readable</b> 22:25 23:2,3,4,5,12			



[reporter - sale]

76:10 92:15 117:4 157:22 158:12 160:1,5 <b>reporter's</b> 5:5 <b>reporting</b> 58:17 <b>represent</b> 109:5 <b>represented</b> 69:9 <b>request</b> 97:16 <b>requests</b> 31:9 <b>require</b> 68:5,6 145:10 148:24 149:1,6 <b>required</b> 122:18 124:1 149:14 <b>requirement</b> 124:15,25 <b>requirements</b> 140:14 <b>research</b> 41:20 61:20 <b>reside</b> 18:2 <b>resize</b> 25:16 <b>resized</b> 25:12 <b>resolution</b> 15:23 <b>resolved</b> 132:23 <b>resource</b> 41:11 54:17 <b>resources</b> 7:9 32:4 65:16 66:8 143:10 144:14 <b>responding</b> 87:22 <b>response</b> 10:12 17:20 33:2 38:16 159:4 <b>responses</b> 12:16 12:18 <b>responsibility</b> 44:11 <b>responsible</b> 19:14 <b>restore</b> 134:14	<b>restricted</b> 136:24 137:12 138:2 <b>result</b> 36:7 129:17 135:24 <b>results</b> 18:8 37:13 129:16,23 136:14 136:22 <b>retain</b> 25:13,14 103:16 <b>retention</b> 57:1,4 <b>retrieve</b> 92:5 <b>return</b> 70:13 81:22 135:9 <b>returned</b> 149:8 <b>returning</b> 38:7 148:25 <b>review</b> 14:4 20:22 20:22,25 21:11,24 27:3,7,7,19,21,24 27:25 28:9,11,12 29:4 44:12 45:2 54:5 56:3,19 68:5 68:7 72:24 73:4 133:18 140:24 142:1,15 154:4,15 <b>reviewed</b> 7:20 133:10 <b>reviews</b> 73:1,24 <b>rich</b> 148:11,15 <b>rid</b> 110:8,9,19 111:9,19,21,25 112:8,12 114:6,6 114:18,25 <b>rids</b> 113:2 <b>right</b> 6:1 8:2 9:3 9:11,19 10:2,17 11:9 12:4 16:21 17:25 19:4 20:7 20:13 21:16 22:4 23:22 24:7,7 27:11,12,15 31:1	31:11,13,15 32:15 32:23 35:16 43:16 43:24 45:7,11,17 45:18,24 47:19 49:1,22 50:2 51:13 55:4,20 56:17 60:7 61:10 61:22 62:21 63:7 63:10 66:1 67:23 69:18 72:11 74:2 80:14,17 82:5,22 86:19 88:3 89:3,5 89:17 92:3 95:23 96:12 97:17 100:21,22 103:12 104:9 106:25 107:3 108:25 110:21 112:14 120:22 121:2,4 125:14,15,25 126:5,25 127:1,18 133:25 137:12 138:16 141:12,14 143:7 145:14 148:13 149:13 154:2,11 <b>ring</b> 3:21 10:10,19 11:9,10 13:1 19:9 19:23 20:1,3 38:13 39:1,4 40:1 40:7,11,15,21,25 48:10,18 57:23 60:12,21 63:12,23 77:19 78:9,12,17 93:20 98:17,25 99:15,22 101:6,18 101:20 104:17 105:1,11 107:6,11 107:16,18,21 110:14,21,25 111:1,5 116:5,9,11	130:16 133:3 135:9 139:8,17 145:22,25 147:19 148:9,14 150:1 151:3,8,15,17 152:5,9,18 153:10 153:24 157:24 158:13 159:2 <b>road</b> 3:17 <b>rohrback</b> 3:9 <b>roi</b> 81:16,17 <b>room</b> 126:14 <b>rooms</b> 6:19 63:17 63:18 <b>rosemarie</b> 3:21 <b>row</b> 44:19 130:7 130:15 145:16,21 <b>rows</b> 43:15 44:2 101:22 129:24 <b>rpr</b> 1:18 160:25 <b>rring</b> 3:23 <b>rule</b> 56:22 <b>rules</b> 13:11,12 122:23,24 123:7 124:23 125:17,22 <b>run</b> 35:18 39:13 41:13 73:23 80:18 81:9,25 112:19 129:11,13,22 130:9 131:13,17 136:21,21 154:5 155:5 <b>running</b> 39:14 41:12 67:13 106:1 <b>runs</b> 73:21,22 75:15
			<b>s</b>
			<b>safeguards</b> 108:4 <b>sake</b> 5:5 27:18 <b>sale</b> 70:17

## [sales - simple]

<b>sales</b> 70:8,9	111:16,17 118:4,8	<b>seen</b> 34:11 70:12	139:15 153:7,13
<b>samples</b> 101:12	118:14 120:6	113:11	160:8
<b>san</b> 3:22	121:13 122:2,8	<b>self</b> 117:16	<b>sets</b> 79:8 87:16
<b>sandeep</b> 4:3	123:8,8,25 124:19	<b>send</b> 95:21 156:23	<b>settings</b> 89:1,2
<b>saw</b> 14:2,4 30:9	124:24 125:1,3	<b>sense</b> 17:19 18:20	124:10
39:16 72:2 99:23	126:18,19	25:2 27:11 31:15	<b>seven</b> 143:12
101:3 102:17	<b>schemas</b> 43:1,5,23	59:24 99:20,20	<b>share</b> 56:13
127:8 132:15	59:8 84:7 85:6	100:21 105:23	149:16
<b>saying</b> 6:18 19:11	114:23 117:14,17	112:10 123:17	<b>shared</b> 73:3 94:10
23:12,12 24:12	117:21,25 154:8	125:24,25 128:24	156:6,18
28:22,23 47:24	<b>schematization</b>	129:6 136:16	<b>sharing</b> 133:1
58:22 74:7 96:24	30:18 121:21	144:1 146:24	<b>sheet</b> 80:19
96:25 98:20	<b>school</b> 57:25 58:1	147:2	<b>shelf</b> 129:20
107:24 110:14	<b>scope</b> 6:8	<b>sent</b> 96:21	<b>shopping</b> 46:23
111:23 135:15	<b>screen</b> 17:21 24:5	<b>sentences</b> 63:22	94:8
147:10,15	59:19	150:15 151:6	<b>short</b> 23:16
<b>says</b> 13:11 19:15	<b>script</b> 36:8 129:20	<b>separate</b> 117:22	<b>shorthand</b> 2:11
33:2 56:20 66:11	<b>scripts</b> 35:6	131:23	9:7,15,23 12:8
89:21,21,23 94:5	<b>scroll</b> 44:7	<b>separateable</b>	160:4,12
128:12 136:17	<b>sdk</b> 64:17,24 66:12	113:9,10	<b>show</b> 23:1 59:17
141:10 145:16	66:15 67:16 79:15	<b>sequence</b> 41:13	101:12 124:8,11
150:17 151:11	<b>sdks</b> 68:24	139:4	124:11,12,18
<b>scale</b> 55:10 108:23	<b>se</b> 157:12	<b>sequential</b> 35:14	<b>shut</b> 41:3
112:25 119:20	<b>search</b> 72:8	102:2	<b>sid</b> 111:19 113:9
<b>scaleable</b> 124:24	134:16	<b>series</b> 12:15 43:9	113:14,23 114:5,7
<b>scales</b> 120:10	<b>seattle</b> 3:11	50:12	114:18
<b>scaling</b> 120:12	<b>second</b> 32:16 46:9	<b>serious</b> 29:8	<b>side</b> 25:21 79:20
<b>scan</b> 130:12	62:9 79:2 86:3	<b>serve</b> 135:19	142:3
<b>scenario</b> 53:12	<b>secret</b> 45:16	<b>servers</b> 126:17,20	<b>sids</b> 113:17
<b>scenarios</b> 111:24	<b>secrets</b> 45:20 46:3	126:20	<b>sign</b> 42:12 68:5
153:5,6,8,14,17,20	59:10	<b>serves</b> 102:16	110:1 138:22
156:23 158:16	<b>security</b> 26:4 59:1	<b>service</b> 66:10	139:21
<b>schedule</b> 41:14	155:8	136:17 137:11	<b>signature</b> 160:23
75:14	<b>see</b> 16:10 17:21	<b>set</b> 6:24 15:6 31:23	<b>significant</b> 67:23
<b>scheduler</b> 35:5,19	46:24 49:15,15	44:19 49:20 50:24	74:3 109:14
36:12	53:2 70:16 83:2,5	58:24 69:6 83:19	<b>similar</b> 15:7 16:14
<b>schema</b> 30:8,13,16	83:7 86:23 94:17	84:6 85:18 87:19	98:1 114:5 118:25
30:21 32:3 43:14	94:18 124:1,3,14	87:23 89:6 97:22	119:20 135:20
43:14,25 44:1	128:4 144:20	99:9 105:15	137:4
45:6 50:3,25 51:4	145:19	122:22,23 129:13	<b>simple</b> 68:4
51:6 95:9 111:15		136:6,19 138:8	133:19

**[simpler - stamp]**

<b>simpler</b> 148:10	<b>somewhat</b> 15:19	<b>spared</b> 132:22	110:11 111:11,14
<b>simply</b> 134:20	<b>soni</b> 88:17	<b>spark</b> 33:15 34:6	112:3 118:4 120:1
<b>single</b> 15:4 17:10	<b>soon</b> 40:17	35:1	120:2,18 123:16
67:12 85:23 91:18	<b>sooner</b> 112:17	<b>speak</b> 6:18 57:12	129:24 130:25
91:24 92:1 101:22	<b>sorry</b> 14:2 24:16	73:21 74:20 81:14	132:6 135:12
106:3,4 108:23	28:19 30:19 34:8	88:13 116:17,23	137:3,5 138:9
109:9,10,17,23	38:14 39:20 40:1	142:25 150:10	139:5,23 140:14
130:12 134:5	48:10 57:21 65:12	154:22 156:8	141:4 146:3 147:2
135:24 149:3	76:11,13 78:22	<b>speaking</b> 19:2	147:23 149:4
156:21	79:24 80:11 92:15	34:2 40:21 41:5	153:1,8
<b>sit</b> 30:19 60:5	99:24 104:17	82:22 103:15	<b>specifically</b> 36:25
92:19	107:6 108:1 109:1	<b>special</b> 1:11 6:2	48:24 57:6 60:1
<b>site</b> 15:3,25 16:1	109:1 110:14	10:10 11:10 19:9	64:20 70:9 71:24
16:17,18	113:21 116:5,22	38:13 60:12 62:10	78:19 80:6 84:6
<b>sits</b> 42:17 53:10	116:25 121:25	71:12 72:25 76:11	85:15 87:14
115:7	133:4 134:17	77:19 93:10 96:22	100:17 114:21
<b>sitting</b> 16:18 30:14	137:16,17 145:20	98:17 101:6	119:17 121:19
141:11	151:3 153:10	106:22 123:18	122:10 123:7
<b>situated</b> 8:24	154:25 155:13	130:16 132:21	135:13 149:21
<b>situations</b> 146:17	158:1	133:3 145:2	155:1
<b>six</b> 138:17 141:8	<b>sort</b> 11:19 18:9	147:19 148:18	<b>specificity</b> 61:19
141:12	54:25 56:21	158:14	<b>specifics</b> 119:18
<b>size</b> 15:22 61:24	103:11 142:14	<b>specialist</b> 132:2	<b>speculate</b> 76:22
106:8 127:15	<b>sorts</b> 23:22 25:6	<b>specialists</b> 101:7	148:21
<b>skimming</b> 16:2	55:4	<b>specialized</b> 18:10	<b>speed</b> 100:23
<b>sleep</b> 14:6	<b>sought</b> 135:1	142:2	<b>spend</b> 70:14 81:14
<b>slightly</b> 15:6	<b>sounded</b> 150:2	<b>specialty</b> 10:22	81:22
<b>slow</b> 37:11 91:3	<b>sounds</b> 41:9 53:21	<b>specific</b> 10:14	<b>spending</b> 81:18
<b>small</b> 15:10 130:4	55:22 84:5 92:10	16:24 18:3 37:10	<b>spent</b> 49:16
<b>smaller</b> 25:14	92:12 102:2 103:9	39:9 41:5 52:23	110:17
<b>social</b> 54:15 64:24	105:25 150:9	53:19 56:12 66:2	<b>spoke</b> 155:23
<b>software</b> 10:22	152:2	68:1 69:17,20	<b>spoken</b> 140:22
34:3 67:11 72:10	<b>source</b> 81:3 89:8	70:11 74:11 81:24	<b>spread</b> 51:24
121:18,19 122:9	109:9,10,17,23	83:25 84:7 88:19	<b>spreadsheet</b> 81:11
122:15	129:12,12	90:21 91:9 93:25	<b>sql</b> 34:7,9,9,10,12
<b>solanki</b> 4:3 157:24	<b>sources</b> 88:1	94:2 95:9,11,15,18	34:16,16,17,17,19
158:1	149:24 150:4	96:14,18,25 97:2	34:19,24 35:1
<b>somebody</b> 22:19	<b>space</b> 84:1,16	98:15 100:13	40:20 41:10,10
53:11 74:15 79:6	120:3	101:25 103:19	<b>ssd</b> 76:25
103:22 120:5	<b>spacer</b> 150:17	104:11,13,14	<b>stamp</b> 149:15
124:21 150:7		105:10 108:11	

## [standard - table]

<b>standard</b> 2:13 34:9,17 97:7,22,24 99:8 139:11,20,22 142:9 <b>standardized</b> 140:17 <b>standpoint</b> 37:18 <b>stands</b> 27:20 <b>start</b> 13:7 14:24 20:16 39:23 94:11 102:16,16 119:18 143:19 145:12 <b>started</b> 58:12,16 64:4 139:6 <b>starting</b> 117:12 <b>state</b> 2:12 62:3 72:8 88:23 160:5 <b>stated</b> 149:11 <b>statement</b> 31:18 31:20 92:8 <b>statements</b> 12:23 <b>states</b> 1:1 2:1 135:5 <b>step</b> 79:2 <b>stick</b> 20:14 44:25 99:4 <b>stopped</b> 58:14 67:9 <b>storage</b> 17:8 33:12 33:14 69:10,21 84:18 91:1,3 103:8 110:4 113:1 115:6 119:9,23 120:9,10,22 121:3 123:14 125:9 127:11 129:25 131:21 132:4,11 133:1 134:5,7,13 134:14,21,23 135:4 136:18,22 138:13	<b>store</b> 18:3,10 36:2 36:17,17 54:18 55:21,24 69:11 70:8,11,16 75:4,13 76:18 79:8 86:8 91:23 100:9 106:10,11 113:23 126:10,13 134:20 135:24 148:22 <b>stored</b> 17:7 18:5,6 18:7 32:19 35:25 36:18 44:19 50:10 50:13 51:14,16 54:10,11,12,12,22 54:23,23,24 55:5,7 55:7,22 69:3,24 71:9 74:13,25 75:4,19 76:8,14 80:22,23,24,24 81:2,19,23 83:23 84:21 85:1,4,4 86:13 92:2 96:12 104:11 105:10,20 105:21 107:25 110:19 126:15,16 127:2 135:5 146:6 146:6,7,11 147:1 147:17 148:11,12 148:21 152:11,14 152:21 153:1 <b>stores</b> 85:8 127:25 128:8 <b>storing</b> 16:8 36:2 76:19,21,21,23 127:5 130:3 156:5 <b>straight</b> 78:23 <b>strategic</b> 81:14 <b>street</b> 3:5,21 <b>strike</b> 29:14 63:21 <b>string</b> 95:20 109:7	<b>strip</b> 26:21 <b>strong</b> 128:13 <b>structure</b> 21:20 85:2 108:15 <b>stuff</b> 30:9 31:14 36:5 132:20 <b>sub</b> 136:14 <b>subject</b> 139:10,21 <b>submissions</b> 133:11 <b>subscribed</b> 160:19 <b>subset</b> 16:22 36:18 78:1,4 92:13 120:14,17,17 130:4,11 <b>subsets</b> 79:23 155:4 <b>substantial</b> 65:15 66:8 <b>sufficient</b> 19:3 <b>suggest</b> 60:13 102:18 <b>suggested</b> 153:15 <b>suite</b> 3:5,11,21 <b>summarize</b> 77:11 <b>summary</b> 70:6 77:18 <b>superset</b> 31:24 <b>support</b> 66:14 67:20 <b>supported</b> 24:9 <b>supporting</b> 137:13 138:3 <b>supposed</b> 42:8 108:11 <b>sure</b> 8:20 14:17 16:13 24:3 28:6 63:23 73:3,5 77:1 79:4 80:1 81:15 85:20 89:3 96:22 98:11 99:16	107:12 108:2 109:11 124:2 139:7 144:6 148:22 149:13 152:20 <b>surface</b> 125:18 <b>surfaces</b> 21:5 <b>suspect</b> 52:8 <b>swear</b> 8:4,9,10,12 11:24 <b>swimming</b> 88:16 88:18 <b>switched</b> 56:15 <b>sworn</b> 9:3,6,11,14 9:19,22 10:2 12:4 12:7 41:1 <b>synonymous</b> 33:4 33:13 <b>system</b> 17:10,19 18:13 25:18 33:12 82:3 85:8 101:11 109:20 113:1,15 121:3,22 146:12 146:24 153:1 154:20 <b>systems</b> 8:22 13:3 17:8,11,14,16 18:22 19:7 30:12 30:13 36:18 37:19 76:5 81:22 92:9 102:20,25 103:1 105:22 106:10,19 106:21 107:5 108:10 117:20 119:2,4,4,6,13,16 119:22 154:21 158:6  <b>t</b>  <b>table</b> 30:20 35:9 35:11,12 37:25 38:9,10,11 43:8
--	--	---	--

## [table - things]

44:22 45:10 50:3 50:24 51:6 75:24 102:15 104:25 105:25 106:4 107:2,3,8,11 108:21,24 109:20 111:14 112:12 115:2 118:7,9,11 118:13,14 120:5,6 121:9,23 130:8,11 130:13 141:11 144:15 145:6 151:21 156:20 <b>tables</b> 39:7 41:9 43:17,18,22,24 44:2 50:4 59:9 60:3 92:13 111:15 111:16,18 112:8,9 114:17 115:3,12 115:19 118:4,8,17 121:9,10,11 130:19,20 131:5 132:3,24,24 133:20 145:3 <b>tabs</b> 116:16 <b>tag</b> 22:21 23:17 <b>tagged</b> 14:7,13 22:23 86:16 113:2 <b>tailor</b> 140:14 <b>tailored</b> 140:8 <b>take</b> 5:5 29:8 40:12 60:19,20,22 63:7 77:10 93:3 93:12 112:25 117:5,6 127:10 130:20 133:17 153:17 <b>taken</b> 2:10 63:25 117:8 120:5 160:7 <b>takes</b> 15:20,21 22:19 123:15	129:1 134:15 145:5 <b>talk</b> 30:20 35:16 41:2 58:8 84:1,16 100:13,16,20 121:17 141:3 152:2 157:17 158:22 <b>talked</b> 34:5 90:10 102:21 108:3 110:4 130:18 147:21 158:3 <b>talking</b> 11:13 13:9 13:25 14:17 38:9 38:18,21 43:23 45:5 46:3 48:25 61:2 62:6,11,16,20 71:24 78:11,16,18 78:19 80:2 82:20 107:9 108:4 110:22 111:24,25 111:25 135:7 146:12 147:7,9 151:19 <b>tangible</b> 123:9 <b>tao</b> 36:23,24 37:7 37:11,13,18,22 52:24,25 54:12 90:19 91:1 92:2,5 92:13,19,24 105:16,21 106:6 106:25,25 108:7 109:18,19 110:2,6 117:12,13,15,17 117:18,21 118:4 118:13,18 119:13 119:17,23 120:1 120:14,17,21 121:4,6,7,8 125:10 125:10,21,21 126:10,13,14,16	126:19,22 127:4 127:17,20,24 128:8 137:6 145:12,13 154:5 <b>tao's</b> 106:8 <b>target</b> 129:12,14 <b>targeted</b> 144:12 <b>targeting</b> 147:7,24 148:3,15,20 150:7 150:10 151:22 152:5,6,10,18,20 <b>task</b> 13:5 27:13 28:2 29:1,21 30:2 35:7,11,13 101:21 101:24 102:5,13 102:15,16,17 <b>tasked</b> 113:14 <b>tasks</b> 28:13 42:18 101:3 102:6,12,19 <b>team</b> 22:5 27:6,19 27:21 28:3,4,6,14 28:15 52:16 53:19 54:16 56:2,3 73:24 77:17,17 91:22 100:17,20 118:25 119:1 142:3 146:19 148:20 150:7,10 150:24 151:1 152:2,24 <b>teams</b> 10:22 142:4 142:7 <b>technical</b> 7:4 21:4 31:4 115:6,8 134:4 <b>technically</b> 38:12 133:20 <b>technique</b> 109:5 <b>tee</b> 19:1 101:16 <b>tel</b> 3:6,12,18,22	<b>tell</b> 5:7 14:7 19:3 19:19 23:9 53:23 53:24 58:6 74:12 98:14,19,20,24 99:2 113:16 148:22 <b>telling</b> 90:16 133:8 <b>temporary</b> 41:17 <b>ten</b> 63:11,12 <b>tends</b> 117:19 <b>terabytes</b> 135:3 <b>term</b> 23:16 151:18 <b>terms</b> 30:13 38:18 41:5 58:13 139:11 139:22 <b>test</b> 5:13 <b>testified</b> 9:8,16,24 12:9 <b>testifying</b> 19:13 160:10 <b>text</b> 36:11 127:3,6 <b>thank</b> 6:22 7:7,22 8:5 10:6 11:10 12:24 20:4 31:3 41:3 93:15 116:9 157:5,5,8,9,15,19 159:2,5 <b>thereof</b> 160:14 <b>thing</b> 14:19 29:18 54:20 55:13 78:13 84:19 132:15 134:17 148:23 151:11 158:13 <b>things</b> 7:10,10 10:17 12:23 16:14 16:24 17:16 18:19 20:23 23:14 25:6 25:24 26:1 28:15 28:16 31:22 44:4 45:7,14,22 50:16 50:17 55:4 79:16
---	---	---	---

**[things - transient]**

87:5 89:15 90:7,9 94:24 97:2 98:22 108:11,13,13 116:13 118:24 131:5 141:14 143:6,7 149:21 <b>think</b> 6:17 7:16,21 10:3 13:23 14:16 14:20 16:9 17:9 19:18 22:14,17 23:24 25:25 26:2 30:11,13 31:19 39:2 44:15 45:21 49:6 51:13 54:13 57:23 58:9 62:5 70:19 72:19 75:19 76:18,20 83:3 85:15 86:22 87:5 87:7,9 90:14,15 93:8,9 97:14 99:11,23 102:20 102:25 103:2,9 104:2 105:3 107:21,24 111:2 113:12 116:14 120:12 121:15 123:12 124:20 126:7 127:22 130:22,24 133:5 133:10 135:15,18 138:5 139:9,19 144:9 148:2 150:1 150:6 152:23 153:3,16 154:8,11 154:12 155:5 157:1,2,24 158:3 <b>thinking</b> 15:9 73:7 76:13 98:9,22 133:5 153:7,19 154:1 156:23	<b>thinks</b> 8:11 54:16 <b>third</b> 3:11 13:15 30:5 38:23,24 39:2,11 47:21 64:11,14,16,17,18 67:20 68:16,24 71:16 72:1 74:5 86:12,15 88:11 99:7 113:18 114:3 114:10,11 115:24 138:6,7 140:20 141:9 146:21 149:17 155:3 156:7 <b>thought</b> 20:11 21:20 47:6 78:7 80:20 113:10 134:25 135:11 138:2 <b>thoughtful</b> 157:3 <b>thousands</b> 126:16 126:17 <b>three</b> 13:19,21 34:4 63:22 100:22 <b>thumbs</b> 63:4 <b>tie</b> 29:2 <b>tied</b> 23:18 95:2 <b>tier</b> 129:15 136:13 <b>time</b> 2:13 5:7,8,11 7:1,3 8:15 11:13 15:17,20,20 18:9 32:13 49:16,18 52:16 53:2,4 56:15 59:17 62:9 63:13 65:16 75:8 76:7 87:13,14 90:24 91:12,18,24 92:1,23 94:20 110:17 117:20 119:10 128:18 129:1 130:9,14	132:16 133:17,17 133:18 134:15,19 135:2,17,25 136:8 136:21 143:2 145:5,12 149:3,15 153:18 154:2,3 155:13 157:10,14 157:14,19 158:16 158:18,20 159:5 160:8 <b>timeline</b> 6:23 54:12,13,14 55:24 <b>timelinedb</b> 18:7 18:11 20:6 51:19 52:1,10,11,11 <b>times</b> 38:5 75:7 <b>today</b> 6:4 12:14 13:7 19:13 60:5 62:16,20 66:7,23 70:22,23 101:14 116:13 142:17,22 154:12,16 157:7 <b>today's</b> 6:3,12 <b>told</b> 99:6,21 104:4 <b>ton</b> 61:22 <b>tool</b> 42:21 46:25 57:21,21,22 58:6 65:5,23 66:6 68:20 69:3 70:3,5 70:20 77:8,12,23 78:1,3,11,12,14,19 78:20,20,24 80:7 80:11,12 81:3,7,8 81:20 82:2,5,11,16 85:12,15 91:21 93:25 95:18 97:3 98:15 114:12 139:10,20 <b>tools</b> 64:21,23 66:7 66:21,23,23 68:17 69:5,7,13 74:10,12	74:24 76:3 77:4 79:12 83:16 94:23 <b>top</b> 33:11 34:4 84:23 97:9 125:8 <b>topic</b> 19:10 110:16 <b>topics</b> 10:14 <b>total</b> 85:13 <b>touch</b> 17:17 39:3 39:12 <b>touching</b> 39:7 130:8 <b>tour</b> 59:23 91:17 92:1 <b>track</b> 5:6,11 29:2 74:5 82:25 102:5 102:11 <b>tracked</b> 55:16,19 <b>trade</b> 45:16,20 46:2 59:10 <b>traditional</b> 30:13 <b>train</b> 134:25 135:11 <b>trained</b> 20:20 28:7 29:5 42:12 108:10 <b>training</b> 42:6 122:17 <b>transaction</b> 128:19 <b>transcribe</b> 63:1 <b>transcribed</b> 160:13 <b>transcript</b> 63:6 157:3 <b>transcription</b> 160:14 <b>transfer</b> 75:3 <b>transferred</b> 103:8 <b>transformation</b> 144:23 <b>transient</b> 23:16,16
---	--	--	---



[translatable - user]

<b>translatable</b> 94:24 95:7 <b>transparency</b> 14:25 78:1 79:19 148:6 <b>travel</b> 23:14 <b>treat</b> 24:4,6 42:8 54:18 <b>treated</b> 158:11 <b>tree</b> 123:5 <b>tremendous</b> 125:16 <b>tried</b> 29:19 49:6 49:20 101:11 <b>trigger</b> 113:3 <b>true</b> 72:13 100:6 <b>truly</b> 67:8 <b>truth</b> 109:9,17,24 <b>try</b> 52:13 64:9 113:11 155:12,13 157:6 <b>trying</b> 14:11 19:18 19:19 24:18 25:21 25:23 29:22 39:5 39:15 41:20 42:15 42:17 43:11 50:1 50:5,8 52:18 53:24 54:21 57:19 70:19 74:8 80:13 82:3,15 85:10 89:7 101:10 120:16 131:11 139:3 140:3,5,6 141:18 146:18,22 148:13 154:2 156:14 157:24 <b>turn</b> 65:20 116:21 158:18 <b>turnaround</b> 158:16	<b>turned</b> 66:9 <b>tv</b> 24:1 <b>tv</b> s 23:21,23,24 <b>twice</b> 101:1 <b>two</b> 7:2 13:19,21 22:17 32:7 71:5 83:14 98:22 135:20 137:8 150:15 151:6 154:8 <b>tying</b> 71:10 <b>type</b> 64:11,14 85:1 87:5 88:4 121:9 127:2 136:7 149:15 <b>types</b> 22:1 45:13 79:16 118:16 138:12 142:5 150:12	64:9 65:24 68:20 70:2,8 81:9 82:4 87:8 96:5 107:1 115:15,18 121:13 123:14,21,23 126:11 128:11 133:5 139:3 144:3 144:6 150:20 154:21 155:9 158:24 <b>understanding</b> 19:20 20:1 31:12 52:11 68:2 100:15 125:12 132:9 <b>understood</b> 32:23 47:14 61:4 68:13 85:21 98:11 137:25 155:23 <b>undertaking</b> 119:20 <b>unified</b> 50:22 <b>unintelligible</b> 96:3 99:5 <b>unique</b> 84:15 <b>unit</b> 33:11 35:7 <b>united</b> 1:1 2:1 96:20 135:5 <b>units</b> 81:12 128:13 <b>universe</b> 136:10 <b>unix</b> 36:9,9 <b>unreadable</b> 23:10 <b>unsupported</b> 65:19 <b>unzip</b> 133:2 <b>update</b> 37:12 67:11 <b>updated</b> 117:20,20 117:21 <b>updates</b> 126:19 <b>upload</b> 22:20 25:13 26:11 70:10	75:9 80:18,21 81:9,10,11 82:3 83:16 123:25 <b>uploaded</b> 16:24 24:25 75:5,8 84:4 86:2 <b>uploading</b> 81:19 82:9 <b>usable</b> 144:24 <b>usage</b> 15:21 49:11 50:19 53:14 142:6 <b>use</b> 7:22 19:16 26:16 33:15 35:2 47:22 51:1,3 52:7 52:21 53:8 65:21 68:2 69:18,19 71:4,18 74:19 75:14 76:16 87:6 90:23 106:24 111:15,16 120:7 122:2,6,19 123:10 123:22 129:22 136:24 137:3 138:2,12 139:1,16 139:21 145:9 146:18,19,20,22 146:25 150:11 152:21 <b>useful</b> 143:25 157:2,19 <b>user</b> 1:6 2:6 13:12 22:16,17,19 26:7 42:20,21,22,24 43:9 44:13 45:25 46:3,4,7,12,13,20 47:3,7,8,12,18 48:20,24 49:15,17 50:23 51:2 53:17 54:1 55:6,21 57:20 59:10 75:5 75:6,18,20,22 76:6
	<b>u</b>		
	<b>uh</b> 82:7 <b>uid</b> 109:9 113:19 113:22,25 114:4 114:17,25 121:7 <b>uids</b> 114:3 <b>ultimately</b> 21:10 23:18 <b>unacceptable</b> 11:3 <b>unclear</b> 48:8 <b>uncommon</b> 36:5 <b>underlie</b> 19:6 <b>underlying</b> 123:14 <b>underneath</b> 125:22 <b>underscores</b> 97:21 <b>undersigned</b> 160:4 <b>understand</b> 14:11 18:1 29:18 32:18 34:8 39:15 42:15 42:17 48:8 49:24 50:1,5,8 57:17		

[user - weaver]

76:15 84:4 85:22 85:23 89:4,15 90:25 91:7,10,12 91:16 95:3 100:11 103:5,16,21 104:7 104:10,21,21,22 104:24,25 105:10 105:14,16,18,19 105:20,24 106:5 106:11,15,21,23 107:1,2,4,7,10,11 107:16,19,20,25 108:18 109:3 110:9,11,19 111:9 111:11,15,18,21 111:25 112:3,12 119:17 120:18 121:10 123:22 124:4,7,17 127:13 129:4 130:10,19 130:21 136:6 141:21 145:18 146:3,8,8 149:17 150:11 156:16 <b>user's</b> 13:14,15,16 45:6 49:12 82:9 <b>users</b> 46:2 49:10 49:13,14 59:12 70:12 77:13 96:4 109:5 156:6 <b>uses</b> 105:5 109:4 127:15,15 136:25 137:14 139:20 <b>usually</b> 8:6,7,8 90:18,21,22 126:18 140:10,11 146:17,22 <b>utility</b> 92:18,21 <b>utilized</b> 142:13 156:18,19	<b>v</b> <b>validate</b> 110:1 112:20 <b>validation</b> 113:4 <b>value</b> 24:12,13,20 26:13 43:12 45:9 92:18,22 98:2 114:18 149:6 <b>values</b> 43:23 44:3 99:4 <b>variables</b> 23:8 <b>varies</b> 88:2 <b>variety</b> 20:23 81:21 <b>various</b> 142:16 <b>vary</b> 58:21 <b>vc</b> 1:6 2:6 <b>vendor</b> 97:1 <b>verbatim</b> 160:11 <b>verify</b> 96:2 98:4 99:8 112:18,20 <b>version</b> 15:3 25:14 34:2 35:2 <b>versions</b> 24:8 34:21 <b>versus</b> 43:24 57:7 <b>videoconference</b> 1:14 3:8,13,19,23 4:5,8 <b>view</b> 94:19 95:22 97:7 122:23 124:5 <b>viewed</b> 51:24 95:12 <b>viewing</b> 96:1 <b>violating</b> 58:13 <b>visited</b> 49:11,17 50:19 54:22 <b>visiting</b> 53:8 79:17 <b>visits</b> 54:23 91:12 <b>voice</b> 90:11	<b>w</b> <b>w</b> 3:10 <b>wait</b> 35:9 61:9 104:17 144:19 <b>walk</b> 54:5 65:8,22 70:16 79:3 91:13 140:19 <b>walked</b> 103:18 131:1 <b>want</b> 5:10 6:21,24 7:7,22 8:4,21 10:11,14,15 11:11 12:20,22,24 19:17 22:7 31:7 36:6 38:20 40:4,6,8,9,9 40:25 48:18 59:12 60:15,25 61:11 63:11 64:14 65:13 65:17 68:6 74:7 78:4 81:16 84:15 84:24 85:20 89:2 93:3 94:5,5 96:22 97:20,21 98:10,18 98:21 99:16 100:12,13,14,16 101:3 103:10,21 105:1 107:6,8,12 108:1 115:16 118:7 120:25 130:4 133:4,12 135:10,16 139:16 141:10,13,14,15 149:12 152:11 155:5 <b>wanted</b> 7:10 11:11 11:14 13:4 31:21 68:8 80:1,10 98:21 111:3 117:2 130:18,22 157:5 157:15	<b>wanting</b> 60:14 <b>wants</b> 32:10 70:8 <b>warehouse</b> 32:20 32:21,24 33:2,22 51:15 52:3 69:14 86:14 106:17 115:5 143:13,21 <b>warehouses</b> 32:25 69:19 86:14 131:10 <b>warehousing</b> 85:2 <b>warm</b> 131:21 134:5,7,22 <b>washington</b> 3:11 <b>watched</b> 49:16 <b>watching</b> 49:16 87:20 <b>way</b> 4:3 7:23 8:21 12:14 13:2,4 15:9 24:19 30:12 37:7 46:14 56:21 58:14 59:2 60:3 61:12 61:15 62:20,20 68:22 74:6 79:3 79:13 81:17 84:15 99:8 103:10 106:20 111:2,12 112:18 113:9 119:15 120:1,8 122:1 125:1,3,6 129:10 130:10 133:19 145:15 <b>ways</b> 83:14 103:9 129:18 135:21 <b>we've</b> 11:12,15 12:11 59:11,14 102:21 110:17 133:11 140:21 142:16 <b>weaver</b> 3:4 8:3 46:18,19 47:18,19
--	--	--	--



[weaver - zoom]

48:4,7,11 60:17 61:3,4 62:3,8 93:8 93:15 115:20,21 116:5,8,10 117:6 126:7 157:1 <b>weaver's</b> 49:23 60:20 <b>web</b> 24:5 53:5 <b>website</b> 46:21 68:8 68:11 70:9 71:21 86:4,4,7 <b>wednesday</b> 1:15 2:13 5:1 <b>week</b> 38:5 <b>weeks</b> 7:2 <b>weigh</b> 116:21 <b>welcome</b> 12:23 31:16 88:9 <b>went</b> 46:21 59:22 91:18 <b>whereof</b> 160:18 <b>widgets</b> 81:10 <b>wife</b> 61:12 <b>windows</b> 65:18,21 65:21,25 <b>wish</b> 86:21 <b>witness</b> 9:6,14,22 12:7 160:18 <b>witnesses</b> 6:16 8:4 8:12 10:13 62:13 160:9 <b>wondering</b> 21:15 22:4 59:7 80:8,23 81:18 <b>words</b> 5:15 23:3 <b>work</b> 24:22 26:18 27:16 28:3,4 30:2 39:10 53:1 67:23 74:2 83:17 137:15 143:3 145:9 158:24	<b>worked</b> 10:25 30:12 38:15 48:13 48:22 100:16 <b>working</b> 15:16 50:11 51:8 52:4 67:9 125:23 143:25 <b>workings</b> 82:19 <b>works</b> 52:12 65:19 69:15 81:15 83:23 91:14 99:10 113:24 135:13 139:4 142:3 144:5 <b>world</b> 28:23 123:20 146:20 156:16 <b>worried</b> 48:17 <b>wrap</b> 156:24 <b>write</b> 35:6 73:2 86:24 100:6,7 102:19 119:4,13 119:16 122:14 124:6 125:21 129:2 146:23 <b>writes</b> 103:4 144:20 <b>writing</b> 17:4 100:7 113:15 122:17 145:10 153:5,20 <b>written</b> 48:2 120:20 129:22 153:14,21 154:4 154:13 <b>wrong</b> 102:3 103:10 137:25 <b>wrote</b> 46:5 100:1,3	140:23 142:1 <b>y</b> <b>yeah</b> 11:17 14:25 15:12 16:6 17:20 18:4,6 20:2,3 22:2 22:2 25:20 26:21 33:10 34:1,15,25 36:14 38:10 39:6 39:24,24 40:19,23 43:13,14 44:21 46:4,5,8,10,14 48:21 49:6 51:10 51:12 57:16 60:8 63:15 70:15 71:20 72:7 73:9 75:1 76:17 78:6,16 80:5,15 81:1,4,21 84:25 86:3,10 92:7 95:10 97:5 100:2,6 102:25 106:16 109:16 110:18 112:14 113:5,6,25 114:2,5 114:23 116:15 122:22 123:19 124:20 125:6 126:2 131:25 134:24 135:2 137:24 142:12 145:22,24 148:18 150:9,19 151:14 151:16 152:7,9,13 152:25 153:25 155:7 157:8 159:6 <b>year</b> 141:12 <b>years</b> 66:4,6 67:6 <b>yesterday</b> 145:4 <b>youtube</b> 139:25	<b>z</b> <b>zarashaw</b> 5:19,20 9:2,5 33:4 34:18 37:8 38:4 52:13 53:18 62:19,23 67:1,4,22 69:1,4 69:16,25 90:20 91:8 92:4,21 93:5 105:14 106:3,22 109:11,18,22,25 111:13 112:4,10 116:19 117:15 118:2,6,11,19,24 119:7,14,19,25 120:15,22 121:8 125:4,8 126:12 127:4,14,20 128:1 128:9,15,25 129:10 130:2 134:17 135:18 136:19 137:1,9,16 138:4 143:16,20 144:6,17,25 146:16 150:6 <b>zip</b> 148:17 151:23 <b>zipped</b> 132:4 133:2 <b>zippy</b> 128:6,6,8,11 <b>zone</b> 87:13,14 <b>zoom</b> 1:14 3:8,13 3:19,23 4:5,8
	<b>x</b> <b>x</b> 35:9,11,12 133:17,18 148:5 <b>xfn</b> 27:13,16,19,22 72:25 73:21,22,23		

# Exhibit O

**UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA  
SAN FRANCISCO DIVISION**

IN RE: FACEBOOK, INC. CONSUMER  
PRIVACY USER PROFILE LITIGATION,

This document relates to:

ALL ACTIONS

**MDL NO. 2843**

CASE NO. 3:18-MD-02843-VC-JSC

HON. VINCE CHHABRIA  
HON. JACQUELINE SCOTT CORLEY  
COURTROOM 4 – 17<sup>TH</sup> FLOOR  
SPECIAL MASTER, DANIEL GARRIE, ESQ.

**ORDER FOLLOWING MARCH 9, 2022  
HEARING REGARDING PLAINTIFFS'  
MOTION TO COMPEL PRODUCTION  
OF PLAINTIFF DATA**

**ORDER FOLLOWING MARCH 9, 2022 HEARING REGARDING PLAINTIFFS' MOTION  
TO COMPEL PRODUCTION OF PLAINTIFF DATA**

No later than April 5, 2022, Facebook is to submit a proposed protocol for producing named Plaintiff user data pursuant to Judge Corley's Discovery Order No. 9 (Dkt. 557) beyond what has been produced to date. No later than April 11, 2022, Plaintiffs are to submit a response to Facebook's proposed protocol. Facebook's proposal should include, if appropriate, the following types of information.

- Data flow diagrams of the systems from which Facebook searches and produced Named Plaintiff Data.
- Functional descriptions and interdependencies for DataSwarm tasks that process Named Plaintiff Data.
- Descriptions of the schemas, tables, columns, and data types for Named Plaintiff Data that is produced.

Facebook's proposal should identify any Facebook systems and the Named Plaintiff Data that it will not produce from and include an explanation as to why it will not produce from such systems (i.e., burden, costs, duplicative, etc.)

The Special Master Garrie suggests that Facebook's protocol include producing the Named Plaintiff Data from the following.

- **The Associations Objects (TAO)** – Facebook is to identify and produce Named Plaintiff Data stored in TAO that is not present in the Named Plaintiff's DYI file and it should include, if applicable, the following: (a) the Objects and Associations in TAO associated with the Named Plaintiffs, but are either partially or not included in the DYI file (i.e. Named Plaintiff Data that is not exported in DYI file but stored in TAO). For example, data about a Page (as defined by Facebook) that the Named Plaintiff interacted with is stored in a MySQL table in TAO called Y, which was not included in the DYI file. Facebook identifies table Y and produces the rows/columns from table Y with the Named Plaintiff Data relating to the page.
- **Hive** -- Facebook is to query the Hive to identify the tables that store Named Plaintiff Data using the identifiers including the following: User ID (UID), Replacement ID (RID) Separable ID (SID), App-Scoped ID (ASID). For each table identified Facebook is to search the tables for the associated Named Plaintiff Data. (i.e., tables mapping user identifiers to ad segmentation data such as US Political spectrum segments). For example, a Facebook engineer writes a process that stores user data, including Named Plaintiff Data, in the Hive in table X and the table has a

column “RID.” Facebook would produce the Named Plaintiff Data that is stored in table X. This effort should exclude any of the analysis done in relation to DataSwarm below.

- **DataSwarm** -- Facebook is to query DataSwarm Tasks to identify Task Definitions which involve Named Plaintiffs Data, using known identifiers such as User ID (UID), Replacement ID (RID) Separable ID (SID), App-Scoped ID (ASID). Facebook will review each of those Task Definitions and then search the sources/destinations identified in the task for the Named Plaintiff Data. See Special Master Hearing Transcript 3/9 p.35 5-13. For example, Facebook queries the DataSwarm Tasks and identifies a Task that uses the UserID to pipe data to Laser and stores that data in table Y. Facebook is to search and produce from table Y all Named Plaintiff Data (columns/row/schemas) that was not included in the DYI file.

### Written Questions

Facebook is also to submit answers to the following questions and requests for documentation to the Special Master on or before April 1, 2022.

#### Hive

- How does ad impression and ad click data for Facebook users get into Hive?
- What tables store ad impression and ad click data for Facebook users?
- What data pipelines<sup>1</sup> are used to analyze ad impression and ad click data for Facebook users?

Where is the final output of these data pipelines stored (e.g. Hive table names, TAO, etc.)?

- Identify a list of Hive tables containing columns that store a UID, RID, SID, ASID, or other means of identifying a Facebook user. The list is to include tables that were active during the relevant time period.
- What is the estimated time and cost to produce data for the Named Plaintiffs from Hive?

#### Ads Interests

- How does Facebook determine ads interests for a user based on what the user views (i.e. are ad interests based on what the user views on Facebook or other Internet activity)?
- Does Facebook track user activity across the Internet using cookies? If so, what cookies does Facebook use? Provide a statement explaining the use of cookies in tracking user activity on and

<sup>1</sup> Data pipeline as used here refers to an end to end data process using data tasks in a dependency chain.

off the Facebook platform to create behavioral data about users. Provide documentation on the use of \_fbp, \_fbc, and DATR cookies. Explain whether the scope of tracking user activity includes on or off platform activity, or both.

- How are ads interests associated with a particular user? Where is ads interests data for individual users stored? Facebook is to describe whether ad interests data can be associated with a specific user via UID, RID, SID, ASID, or other means and whether it is included in the DYI file.
- What is the estimated time and cost to produce ad interests data for the Named Plaintiffs?

#### Contracts

- Is the data referenced in the contracts with Netflix, Microsoft, and/or YouTube that Facebook provided to Special Master Garrie for *in camera* review in connection with the Named Plaintiff Data hearings included in the DYI file? If not, Facebook is to specify what data is not included and where such data stored.

#### **Scenarios**

No later than April 4, 2022, Facebook is to submit documentation sufficient to describe the data collected both on and off platform or provided by Third Parties in the following scenarios and provide written responses to the questions below.

- Exhibit A to Plaintiff's Questions re: Data Collection and Use indicates that Facebook used predictive algorithms to generate five political segments for Facebook users (Very Liberal, Liberal, Moderate, Conservative, and Very Conservative) based on demographic, psychographic, and behavioral signals from Facebook user data.
  - What are the inputs into these algorithms (i.e. what are the demographic, psychographic, and behavioral signals used to generate the political segments)? Are these inputs provided by users or derived by Facebook?
  - How are the psychographic signals computed (e.g. how is the psychographic signal "High Dollar Religious Donor" determined)?

- Is information regarding identifiable ethnic affinities provided by users or derived by Facebook? How is ethnic affinity derived?
- Where is political segmentation data for Facebook users stored?
- Is political segmentation determined for a Facebook user as part of a data process that runs on a regularly scheduled basis or evaluated in real time when an ad is served?
- Is political segmentation associated with a Facebook user if possible (i.e. via UID, RID, SID, ASID, or another identifier that can be mapped to a user)? If so, explain how the political segmentation is associated with a Facebook user.
- Is an individual's assigned political segment part of the DYI file?

IT IS SO ORDERED.

Monday, June 27, 2022

---

Daniel Garrie  
Discovery Special Master

# Exhibit P



GIBSON DUNN

Gibson, Dunn & Crutcher LLP

555 Mission Street  
San Francisco, CA 94105-0921  
Tel 415.393.8200  
www.gibsondunn.com

Rosemarie T. Ring  
Direct: +1 415.393.8247  
Fax: +1 415.801.7358  
RRing@gibsondunn.com

HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

April 18, 2022

VIA JAMS ACCESS

Special Master Daniel B. Garrie  
DGarrie@jamsadr.com

Re: *In re Facebook Consumer Privacy User Profile Litigation*, JAMS Ref No.  
1200058674

Dear Special Master Garrie,

Pursuant to the Special Master's Order Following March 9, 2022 Hearing Regarding Plaintiffs' Motion to Compel Production of Plaintiff Data, dated March 22, 2022, Facebook submits this proposal to assist in determining "what, if any" additional named plaintiff data should be produced, as directed in Judge Corley's order dated January 12, 2022 (Dkt. No. 807), and respond to "Scenario" questions.

**Facebook's Proposal For Producing Additional Named Plaintiff Data**

As Judge Corley explained in her most recent order on named plaintiff data, issued on January 12, 2022, the Special Master is working with the parties to determine "what, if any" additional named plaintiff data should be produced consistent with Rule 26. Facebook continues to believe that the named plaintiff data that has already been produced, which includes data in all three categories of "discoverable user data" identified in Judge Corley's Discovery Order No. 9, satisfies its obligations under Rule 26. The additional data Plaintiffs seek is neither relevant to their claims nor proportional to the needs of this case because it was not shared or made accessible to third parties.

That said, in an effort to resolve this issue, and consistent with Judge Corley's Discovery Order Nos. 11 and 12 allowing Plaintiffs to test Facebook's position on whether there is additional named plaintiff data that was shared or made accessible to third parties, Facebook proposes the protocol set forth below for producing additional named plaintiff data.

As explained in Facebook's March 7, 2022 submission to the Special Master, at a high level, user data is stored in production systems and the data warehouse. For production systems we look to TAO, and for the data warehouse we look to Hive. Accordingly, we propose to produce additional data from these systems as explained below.

# GIBSON DUNN

Special Master Daniel B. Garrie  
April 18, 2022  
Page 2

**TAO.** As also explained in the March 7, 2022 submission, TAO is a distributed data store for the social graph. There is an object and associations to that object for all users. *See* Ex. A, Internal Facebook Wiki Regarding TAO Core Concepts. For each named plaintiff, Facebook will produce the user objects and associations to those objects.

**Hive.** As explained in prior submissions to the Special Master, Hive is Facebook’s data warehouse. *See* Ex. B, Declaration of Mengée Ji In Support of Facebook Inc.’s Motion For Reconsideration of Special Master’s Order Regarding Named Plaintiff Data, ¶ 11. Hive is not structured to support centralized searches for individual user data. *Id.* ¶ 16. Instead, it is organized roughly based on product, team, or use case. *Id.* ¶ 17. Searching for specific types of data typically requires seeking the advice of product teams or other subject matter experts to identify potentially relevant tables. *Id.* ¶ 18.

Once potentially relevant tables are identified, they must be individually assessed to determine whether they contain user data and then searched for data related to the named plaintiffs. *Id.* These searches are subject to many technical and policy limitations which make searching for individual user data unduly burdensome, including that (1) many Hive tables do not associate data with individual users (*id.* ¶ 20); (2) many Hive tables are stored in “partitions” or segments, and many of those tables are partitioned by day, reflecting the time when the data was logged or processed, such that each day of data is its own segment and must be independently searched (*id.* ¶ 21); (3) data in many Hive tables is machine readable (*id.* ¶ 22); and (4) data in many Hive tables are in cold storage and would have to be restored in order to be searched and analyzed, *see* Ex. C, Declaration of Mengée Ji In Support of Facebook Inc.’s Motion For a Protective Order Against Production of API Call Logs, ¶¶ 17-18; Special Master’s Nov. 8, 2021 Order Re: Facebook’s Motion For Protective Order Against Production of API Call Logs, ¶ 15 (“Special Master Garrie finds that the data in the Mobile Table and Web Table is not reasonably accessible because it is not readily usable in its ‘cold storage’ state and must be restored to ‘warm storage’ in order to be searched and analyzed (i.e. usable).”).

For these reasons, among others, as reported in its April 11, 2022 submission, Facebook estimates that it would take approximately 62 billion CPU seconds or around 720,000 CPU days of computational cost to search across the entire Hive data warehouse and extract all data about a single user.

Once the searches have been conducted, and data is returned, the data must be manually reviewed to, among other things, confirm that the data is associated with a particular user and does not include personal data of other users (e.g., user blocking another user) or trade secrets, and does not create system integrity or security concerns.

# GIBSON DUNN

Special Master Daniel B. Garrie  
April 18, 2022  
Page 3

In light of the above technical limitations and burdens, Facebook proposes to search for and produce the following named plaintiff data from Hive to the extent it exists:

(1) specific types of data requested or referenced by Plaintiffs in challenging Facebook's production of named plaintiff data; (2) a sample of data from Hive tables with user identifiers included in Exhibit B to the April 11, 2022 submission; and (3) the only data identified in Exhibit C to the April 11, 2022 submission that has not already been produced to Plaintiffs in the DYI file or otherwise: set permissions (audience controls on a post).

## 1. Specific Types of Data

- *Any remaining off-platform activity:* DYI includes off-platform activity. Facebook will search for and produce any underlying raw log-level data for off-platform activity provided to Facebook by a third party associated with the named plaintiffs.
- *Any remaining ad interests:* DYI includes ad interests. Facebook will search for and produce any underlying raw log-level data associated with the named plaintiffs.
- *Any remaining ad click data:* DYI includes ad click data. Facebook will search for and produce any underlying raw log-level ad clicks data associated with the named plaintiffs.
- *Ad impressions data:* Facebook will search for and produce ad impressions data associated with the named plaintiffs.
- *Any remaining custom audience data:* DYI identifies third parties who have created custom audiences associated with a user. Facebook will search for and produce any more granular information about custom audiences associated with the named plaintiffs, including custom audience type, whether it was used to deliver ads, and when.

## 2. Sampling from Hive tables identified in Exhibit A of April 11, 2022 submission

Given that data in Hive tables are not shared or made accessible to third parties, the significant burden of searching for and producing individual user data from Hive as explained in prior submissions, and that Facebook has already produced nearly a million pages of the named plaintiff data from the DYI system, Hive data is not relevant or proportional to the needs of this case. Facebook nevertheless understands that, as Judge Corley found in Discovery Order Nos. 11 and 12, Plaintiffs are allowed to test Facebook's positions on "sharing and accessibility" by obtaining discovery into what data exists and how it is used. To that end, and consistent with the requirements of relevance and proportionality under Rule 26, Facebook proposes a sampling process.

# GIBSON DUNN

Special Master Daniel B. Garrie  
 April 18, 2022  
 Page 4

In addition to the above specific categories of data, Facebook will produce a sample of data in warm storage from Hive tables identified in Exhibit A of the April 11, 2022 submission. Specifically, Facebook will produce data from 200 tables in Exhibit A to be jointly selected by Facebook and Plaintiffs. Facebook will randomly select 100 tables and Plaintiffs will select 100 tables.

### 3. User data identified in Exhibit B of the April 11, 2022 submission and not in DYI

As shown in Exhibit C of Facebook's April 11, 2022 submission to the Special Master, only two types of user data in the contracts submitted to the Special Master for in camera review are not included in the DYI file: (1) privacy settings, and (2) set permissions (audience controls on a post). (1) has already been produced. Facebook will agree to produce (2) from TAO.

With respect to other Facebook systems, not including TAO and Hive, there are seven systems<sup>1</sup> that did not fall into one of the categories identified by the Special Master in Hearing Order Regarding Plaintiffs' Motion To Compel Production of Plaintiff Data, dated February 21, 2022: Ads Raw Storage, Manifold, orderdb (MySQL), Payments (MySQL), Tally, UDB, XDB, and ZippyDB. Producing data from any of these systems would be extremely burdensome and is not relevant or proportional to the needs of this case.

- *Ads Raw Storage and Manifold*: These systems are BLOB and binary storage and are not searchable by user identifiers.
- *Tally*: This system is a real-time aggregation counter and is not searchable by user identifiers.
- *orderdb (MySQL), Payments (MySQL), XDB, and Zippy DB*. These systems contain unstructured data and cannot be searched by user identifiers, except for tables using Ent or Node schemas for which there is a structured deletion plan. Searching other tables for named plaintiff data would require a full text search of all tables that would take many months to complete and generate a significant number of false positives requiring manual review of all data returned.

With respect to Dataswarm, the Special Master suggested that Facebook "query Tasks to identify Task Definitions which involve named plaintiff data" using identifiers,


---

<sup>1</sup> Callisto should have been put in Category (6) because the data is in DYI. UDB should also have been put in Category (6) because it is duplicative of TAO. MySQL describes the general use of MySQL databases and should not be a separate entry. For example, orderdb (MySQL) and Payments (MySQL) are specific uses of MySQL.

# GIBSON DUNN

Special Master Daniel B. Garrie  
 April 18, 2022  
 Page 5

review those Tasks Definitions, and then search for named plaintiff data in the sources/destinations identified. We have investigated the Special Master's suggestion and did not include it in our proposed protocol because it is not feasible. Dataswarm is a collection of data processing operations coded in Python. Task Definitions are the Python code and Tasks are the running of that code. Searching Python code for user identifiers would be done through a syntax-only search which cannot understand the semantic meaning of the data being processed, i.e., whether a column name being processed involves a user identifier. As a result, additional code would have to be developed to understand whether/how identifiers are involved in a particular Task Definition.



## GIBSON DUNN

Special Master Daniel B. Garrie  
April 18, 2022  
Page 6

### Scenarios

No later than April 4, 2022, Facebook is to submit documentation sufficient to describe the data collected both on and off platform or provided by Third Parties in the following scenarios and provide written responses to the questions below: “Exhibit A to Plaintiff’s Questions re: Data Collection and Use indicates that Facebook used predictive algorithms to generate five political segments for Facebook users (Very Liberal, Liberal, Moderate, Conservative, and Very Conservative) based on demographic, psychographic, and behavioral signals from Facebook user data.”

1. What are the inputs into these algorithms (i.e. what are the demographic, psychographic, and behavioral signals used to generate the political segments)? Are these inputs provided by users or derived by Facebook?

**Answer:** As an initial matter, Very Liberal, Liberal, Moderate, Conservative, and Very Conservative are ad interests (“political ad interests”), which are distinct from political segmentation “based on demographic, psychographic, and behavioral signals from Facebook user data” (“political segments”). Political segments are political ad interests that are divided into political segments.

***Political ad targeting segments*** were deprecated for new campaigns in January 2022 and for existing campaigns in March 2022. Based on a reasonable investigation to date, these were ad targeting segments derived from the following on-platform activity (user-provided data and observed data):

[REDACTED]

***Political segments*** were deprecated for new campaigns in October 2017 and for existing campaigns in January 2018. Based on a reasonable investigation to date, these segments were created by

[REDACTED]

## GIBSON DUNN

Special Master Daniel B. Garrie  
April 18, 2022  
Page 7

2. How are the psychographic signals computed (e.g. how is the psychographic signal “High Dollar Religious Donor” determined)?

**Answer:** Based on a reasonable investigation to date, we believe “High Dollar Religious Donor” was an ad targeting option based on partner categories. We do not know how this partner category was created. Partner categories were deprecated in October 2018.

3. Is information regarding identifiable ethnic affinities provided by users or derived by Facebook? How is ethnic affinity derived?

**Answer:** Ethnic affinities, also referred to as multicultural affinities, were ad interests created by Facebook based on a user’s on-platform activity indicating an interest in content relating to certain communities. Ethnic affinity ad interests were deprecated in August 2020.

4. Where is political segmentation data for Facebook users stored?

**Answer:** Political segments were deprecated in October 2017 for new campaigns. Existing campaigns using political segments ended on January 1, 2018. Based on a reasonable investigation, we believe that political segments were deleted in January 2018.

5. Is political segmentation determined for a Facebook user as part of a data process that runs on a regularly scheduled basis or evaluated in real time when an ad is served?

**Answer:** Based on a reasonable investigation to date, we believe political segments were created manually by [REDACTED]

6. Is political segmentation associated with a Facebook user if possible (i.e. via UID, RID, SID, ASID, or another identifier that can be mapped to a user)? If so, explain how the political segmentation is associated with a Facebook user.

**Answer:** Based on a reasonable investigation to date, we believe political segments were associated with users by UIDs.

7. Is an individual’s assigned political segment part of the DYI file?

**Answer:** Based on a reasonable investigation to date, we do not believe political segments were included in the DYI file.

GIBSON DUNN

Special Master Daniel B. Garrie  
April 18, 2022  
Page 8

Sincerely,

A handwritten signature in cursive script that reads "Rosemarie Ring". The signature is written in black ink and is positioned above the printed name.

Rosemarie T. Ring



# Exhibit Q



DESIGNATED HIGHLY CONFIDENTIAL—ATTORNEYS' EYES ONLY BY FACEBOOK

April 29, 2022

**VIA JAMS ACCESS**

Special Master Daniel B. Garrie  
[DGarrie@jamsadr.com](mailto:DGarrie@jamsadr.com)

Re: *In re Facebook, Inc., Consumer Privacy User Profile Litig.*,  
JAMS Ref No. 1200058674, Case No. 3:18-md-02843-VC (N.D. Cal.)

Dear Special Master Garrie:

Pursuant to Order Following March 9, 2022 Hearing Regarding Plaintiffs' Motion To Compel Production of Plaintiff Data dated March 22, 2022, Plaintiffs submit this response to Facebook's Proposal for Producing Additional Named Plaintiff Data, submitted April 18, 2022. For ease of reference, Plaintiffs have pasted Facebook's proposal, "highlighted" in gray, into this document and provide responses to its proposal inline. To summarize, Facebook should be ordered to produce the following data.

From **TAO**, in addition to what Facebook offers to produce—for each Named Plaintiff, the user objects and associations to those objects—it should also be ordered to produce the TAO schema and provide answers to questions about what information is "currently maintained" in TAO.

From **Hive**, for each Named Plaintiff: First, Facebook should produce the types of data it identified in its April 18 submission and the additional types of data identified in this response, whether included in the DYI file or not. However, the production should not be limited to data that is associated with a Named Plaintiff. Rather, it should include data that can be associated with a Named Plaintiff. Second, Facebook should produce the schema and an excerpt comprising the top five rows of each of the Hive tables listed in Exhibit B of its April 11 submission. Plaintiffs propose a process for identifying the tables Facebook should be required to search for data that can be associated with a Named Plaintiff.

Finally, for both **TAO** and **Hive**, based on Facebook's response regarding the time period of the data that it can produce from these sources, it should be required to identify the efforts it made to maintain data that could be associated with the Named Plaintiffs.

**FACEBOOK'S PROPOSAL FOR PRODUCING ADDITIONAL NAMED PLAINTIFF DATA**

As Judge Corley explained in her most recent order on named plaintiff data, issued on January 12, 2022, the Special Master is working with the parties to determine "what, if any" additional named plaintiff data should be produced consistent with Rule 26. Facebook continues

Special Master Danie B. Garrie  
April 29, 2022  
Page 2

KELLER ROHRBACK L.L.P.  
BLEICHMAR FONTI & AULD LLP

to believe that the named plaintiff data that has already been produced, which includes data in all three categories of “discoverable user data” identified in Judge Corley’s Discovery Order No. 9, satisfies its obligations under Rule 26. The additional data Plaintiffs seek is neither relevant to their claims nor proportional to the needs of this case because it was not shared or made accessible to third parties.

That said, in an effort to resolve this issue, and consistent with Judge Corley’s Discovery Order Nos. 11 and 12 allowing Plaintiffs to test Facebook’s position on whether there is additional named plaintiff data that was shared or made accessible to third parties, Facebook proposes the protocol set forth below for producing additional named plaintiff data.

As explained in Facebook’s March 7, 2022 submission to the Special Master, at a high level, user data is stored in production systems and the data warehouse. For production systems we look to TAO, and for the data warehouse we look to Hive. Accordingly, we propose to produce additional data from these systems as explained below.

### **Plaintiffs’ Response to Facebook’s Introduction:**

This is the first time Facebook has conceded that Judge Corley’s discovery orders required the production of any Named Plaintiffs’ data that Facebook contends was not shared with or made accessible to third parties. Before now, Facebook has consistently taken the position that Judge Corley’s discovery orders did not require the production of any Named Plaintiffs’ data that Facebook asserts was not shared with or made accessible to third parties. This about-face is welcome and long overdue.

However, Facebook’s proposed protocol and answers to the Special Master’s inquiry generates more questions than answers. That may be because, as an internal document leaked this week reveals, Facebook does not know where Named Plaintiff data is located (or, indeed, to whom it was disclosed). See Lorenzo Franceschi-Bicchierai, [\*Facebook Doesn’t Know What It Does With Your Data, Or Where It Goes: Leaked Document\*](#), Motherboard (Apr. 26, 2022).

According to the leaked document, Facebook “can’t enumerate all the data we have - where it is; where it goes; how it’s used”—which raises the question, “how can we make commitments about it to the outside world?” The document illustrates the challenge to Facebook using an analogy. “Imagine you hold a bottle of ink in your hand. This bottle of ink is a mixture of all kinds of user data (3PD, 1PD, SCD, Europe, etc.)<sup>1</sup> You pour that ink into a lake of water (our open data systems; our open culture) . . . and it flows . . . everywhere. How do you put that ink back in the bottle? How do you organize it again, such that it only flows to the allowed places in the lake?”

Based on this newly disclosed information, it is possible that the reason Facebook’s efforts to determine what data it has that can be associated with the Named Plaintiffs have been so incomplete, unfocused, and confusing is that Facebook does not and cannot know. If it is not

---

<sup>1</sup> According to the article, “3PD” is third-party data, “1PD” is first-party data, and “SCD” is sensitive categories data.

Special Master Danie B. Garrie  
April 29, 2022  
Page 3

KELLER ROHRBACK L.L.P.  
BLEICHMAR FONTI & AULD LLP

able to identify all data that can be associated with the eight Named Plaintiffs, Facebook should clearly state that fact.

**TAO.** As also explained in the March 7, 2022 submission, TAO is a distributed data store for the social graph. There is an object and associations to that object for all users. *See* Ex. A, Internal Facebook Wiki Regarding TAO Core Concepts. For each named plaintiff, Facebook will produce the user objects and associations to those objects.

### **Plaintiffs' Response to TAO:**

In addition to Facebook's proposal, it should also be ordered to produce the TAO schema referenced on page 2 of Exhibit A in its April 18 submission. Though the reference to the TAO schema exhibit includes a hyperlink, the web page is password protected and inaccessible to Plaintiffs.

In response to Plaintiffs' question about the time range captured by this category of data, Facebook stated that it will search for and produce user objects and associations to those objects "that are currently maintained in TAO which cover the life of the account." Ex. 1. That response raises two more questions. First, what is "currently maintained in TAO"? In other words, for the Named Plaintiffs who are still active on Facebook, are there objects and associations that were once in TAO but are no longer? If so, what has Facebook done to maintain the Named Plaintiffs' objects and associations since the onset of this litigation? Second, what efforts did Facebook make to maintain objects and associations in TAO for the Named Plaintiff who deleted her Facebook account?

Plaintiffs ask that Facebook answer the questions above by May 6, 2022, and that Facebook be ordered to produce the user objects and associations to those objects for each Named Plaintiff, along with the TAO schema, by May 13, 2022.

**Hive.** As explained in prior submissions to the Special Master, Hive is Facebook's data warehouse. *See* Ex. B, Declaration of Mengee Ji In Support of Facebook Inc.'s Motion For Reconsideration of Special Master's Order Regarding Named Plaintiff Data, ¶ 11. Hive is not structured to support centralized searches for individual user data. *Id.* ¶ 16. Instead, it is organized roughly based on product, team, or use case. *Id.* ¶ 17. Searching for specific types of data typically requires seeking the advice of product teams or other subject matter experts to identify potentially relevant tables. *Id.* ¶ 18.

Once potentially relevant tables are identified, they must be individually assessed to determine whether they contain user data and then searched for data related to the named plaintiffs. *Id.* These searches are subject to many technical and policy limitations which make searching for individual user data unduly burdensome, including that (1) many Hive tables do not associate data with individual users (*id.* ¶ 20); (2) many Hive tables are stored in "partitions" or segments, and many of those tables are partitioned by day, reflecting the time when the data was logged or processed, such that each day of data is its own segment and must be independently searched (*id.* ¶ 21); (3) data in many Hive tables is machine readable (*id.* ¶ 22); and (4) data in many Hive tables are in cold storage and would have to be restored in order to be

Special Master Danie B. Garrie  
April 29, 2022  
Page 4

**KELLER ROHRBACK L.L.P.**  
**BLEICHMAR FONTI & AULD LLP**

searched and analyzed, *see* Ex. C, Declaration of Menggee Ji In Support of Facebook Inc.’s Motion For a Protective Order Against Production of API Call Logs, ¶¶ 17-18; Special Master’s Nov. 8, 2021 Order Re: Facebook’s Motion For Protective Order Against Production of API Call Logs, ¶ 15 (“Special Master Garrie finds that the data in the Mobile Table and Web Table is not reasonably accessible because it is not readily usable in its ‘cold storage’ state and must be restored to ‘warm storage’ in order to be searched and analyzed (i.e. usable).”).

For these reasons, among others, as reported in its April 11, 2022 submission, Facebook estimates that it would take approximately 62 billion CPU seconds or around 720,000 CPU days of computational cost to search across the entire Hive data warehouse and extract all data about a single user.

Once the searches have been conducted, and data is returned, the data must be manually reviewed to, among other things, confirm that the data is associated with a particular user and does not include personal data of other users (e.g., user blocking another user) or trade secrets, and does not create system integrity or security concerns.

In light of the above technical limitations and burdens, Facebook proposes to search for and produce the following named plaintiff data from Hive to the extent it exists: (1) specific types of data requested or referenced by Plaintiffs in challenging Facebook’s production of named plaintiff data; (2) a sample of data from Hive tables with user identifiers included in Exhibit B to the April 11, 2022 submission; and (3) the only data identified in Exhibit C to the April 11, 2022 submission that has not already been produced to Plaintiffs in the DYI file or otherwise: set permissions (audience controls on a post).

### **Plaintiffs’ Response to Facebook’s Hive Introduction:**

Because Facebook elaborates on its proposal to produce three types of Named Plaintiff data from Hive in the following sections, Plaintiffs will respond to its proposal regarding those three types of data inline. Plaintiffs respond here to Facebook’s assertion that it will take 720,000 CPU days of computational cost to search across the entire Hive data warehouse and extract all data about a single user.

There is no doubt that there is some burden associated with searching across Hive. But Facebook’s repeated assertion that of the 720,000 CPU days of computational cost required does not provide sufficient information about the actual burden it faces for two reasons.

First, Facebook is not being asked to search across the entire Hive data warehouse. It only needs to search tables that contain data that can be associated with individual users during the relevant time period for this action.

Second, the mention of the number of CPU days of computational cost is almost entirely meaningless without knowing the amount of computing power in Hive. It is reasonable to assume that Facebook, one of the largest technology companies in the world, has available at least tens of thousands of CPUs—if not hundreds of thousands or even more.

Facebook's response also prompts additional concerns. Plaintiffs are concerned that Facebook does not identify the relevant tables that are in cold storage and when they were placed in cold storage. To the extent that any such table were placed in cold storage during the pendency of this action, Facebook should not be permitted to elide its discovery obligations regarding Named Plaintiff data due to its failure to suspend archiving. And Plaintiffs are concerned about Facebook's repeated framing of what it will produce. Facebook's proposal is limited to data that *is* associated with a Named Plaintiff. It should be ordered to produce data that *can be* associated with a Named Plaintiff. In other words, if Hive contains data that is not directly associated with the Named Plaintiff in a table, Facebook should still produce the data if it is possible for Facebook to associate that data with a Named Plaintiff (e.g., by reference to another table).

### 1. **Specific Types of Data**

- *Any remaining off-platform activity*: DYI includes off-platform activity. Facebook will search for and produce any underlying raw log-level data for off-platform activity provided to Facebook by a third party associated with the named plaintiffs.
- *Any remaining ad interests*: DYI includes ad interests. Facebook will search for and produce any underlying raw log-level data associated with the named plaintiffs.
- *Any remaining ad click data*: DYI includes ad click data. Facebook will search for and produce any underlying raw log-level ad clicks data associated with the named plaintiffs.
- *Ad impressions data*: Facebook will search for and produce ad impressions data associated with the named plaintiffs.
- *Any remaining custom audience data*: DYI identifies third parties who have created custom audiences associated with a user. Facebook will search for and produce any more granular information about custom audiences associated with the named plaintiffs, including custom audience type, whether it was used to deliver ads, and when.

### **Plaintiffs' Response to "Specific Types of Data" from Hive:**

Plaintiffs have several concerns regarding Facebook's proposal to produce the above "specific types of data requested or referenced by Plaintiffs in challenging Facebook's production of named plaintiff data."

First, Facebook should not limit its production of these specific types of data based on the specific language Plaintiffs used to describe them. For example, Facebook should not limit its production of ad impressions data to data associated with a specific table, column header, or cell containing the phrase "ad impressions." Unlike Facebook, Plaintiffs do not have access to Hive. Therefore, the "specific types of data requested or reference by Plaintiffs" should be construed broadly as providing descriptive information.

Special Master Danie B. Garrie  
 April 29, 2022  
 Page 6

**KELLER ROHRBACK L.L.P.**  
**BLEICHMAR FONTI & AULD LLP**

Second, Facebook should not limit its production of these specific types of data based on related data included in the DYI file. For example, Facebook states that ad interests are included in the DYI file and agrees to search for raw log-level ad interest data associated with the Named Plaintiffs. But, as discussed below, Facebook also disclosed it does not believe one type of ad interest the Special Master inquired about—political segments—are included in the DYI files. It is not clear whether any other ad interests may not be included in the DYI file. In addition to the raw log-level data, Facebook should also search for and produce *all* ad interests that can be associated with Named Plaintiffs (not just those already associated with Named Plaintiffs), whether included in the DYI file or not. This example should not be understood to indicate Plaintiffs’ interest in ad interests only. Rather, it applies equally to each of the specific types of data identified. For each, it should produce all data capable of being associated with the Named Plaintiffs, whether included in the DYI file or not.

Third, Plaintiffs also request that Facebook produce these additional specific types of data:

- i. Ad profile data for each of the Named Plaintiffs. Aldo King testified Facebook created ad profiles for each user, but also testified that they were uncertain about the scope of data that was incorporated. *See, e.g.,* Ex. 2, King Dep. Tr. at 68:14-69:7 (unsure whether visits to third-party websites were incorporated in ad profiles); *id.* at 83:9-20 (unsure whether transaction data provided by advertisers to Facebook were incorporated in ad profiles). Facebook’s submission also states that a specific type of ad profile (political segments) was based in part on “recent pixel events.” These types of data do not appear to have been included in the DYI files.
- ii. The tables containing data that stores, for the entirety of their time on Facebook, the list of apps authorized or installed, and the period of time during which each such app was authorized or installed, by (a) the Named Plaintiffs and (b) each of the Named Plaintiff’s Facebook friends (by FBID and username). That information should enable Plaintiffs to use other produced documents, such as the Method Table and the Capabilities Tool, to determine whether any of the Named Plaintiffs’ information may have been made available to a third party through friends\_\* or related permissions.

Fourth, to ensure that there is no miscommunication, Facebook should identify in writing the specific tables identified that include these types of information and how it determined those specific tables (and not others) had these types of information.

Fifth, for each of the tables containing this specific information, Facebook should produce data flow diagrams and schemas. *See* March 22, 2022 Order at 1.

Sixth, in response to Facebook’s submission, Plaintiffs asked about the time range of the data Facebook proposes to produce for each of the above categories. In response, Facebook indicated it will produce any *remaining* off-platform activity up to the last two years, any remaining ad interests for “at least 90 days,” and “any *remaining* custom audience data.” Ex. 1 (emphasis added). Those responses are concerning. The Named Plaintiffs joined this litigation more than 2 years ago, and Plaintiffs are concerned about the apparent insufficiency of



Special Master Danie B. Garrie  
April 29, 2022  
Page 7

KELLER ROHRBACK L.L.P.  
BLEICHMAR FONTI & AULD LLP

Facebook's efforts to preserve this information. From Facebook's response of "up to two years," it is unclear whether there is some off-platform activity that is not available for even the two-year period. Further, though Facebook's responses regarding the other specific categories are less troubling on the surface—it states it will produce responsive information from 2015—it is unclear whether Facebook could have produced responsive information from earlier time periods if it had made an earlier effort to preserve the information.

By May 6, 2022, Facebook should be ordered to identify the efforts it has undertaken to preserve the specific types of data identified by Facebook and Plaintiffs in this section that can be associated with the Named Plaintiffs. Facebook should be ordered to produce the above-identified information by May 27, 2022.

## **2. Sampling from Hive tables identified in Exhibit A of April 11, 2022 submission**

Given that data in Hive tables are not shared or made accessible to third parties, the significant burden of searching for and producing individual user data from Hive as explained in prior submissions, and that Facebook has already produced nearly a million pages of the named plaintiff data from the DYI system, Hive data is not relevant or proportional to the needs of this case. Facebook nevertheless understands that, as Judge Corley found in Discovery Order Nos. 11 and 12, Plaintiffs are allowed to test Facebook's positions on "sharing and accessibility" by obtaining discovery into what data exists and how it is used. To that end, and consistent with the requirements of relevance and proportionality under Rule 26, Facebook proposes a sampling process.

In addition to the above specific categories of data, Facebook will produce a sample of data in warm storage from Hive tables identified in Exhibit A of the April 11, 2022 submission. Specifically, Facebook will produce data from 200 tables in Exhibit A to be jointly selected by Facebook and Plaintiffs. Facebook will randomly select 100 tables and Plaintiffs will select 100 tables.

### **Plaintiffs' Response to Sampling from Hive Tables:**

As an initial matter, Plaintiffs believe Facebook intended to reference Exhibit B of its April 11 submission, not Exhibit A. The April 11 submission describes Exhibit A as "a list of data pipelines used to analyze ad impression and ad click data for which we believe output i[s] stored in Hive," and Exhibit B as "a list of Hive tables that" may "store user identifiers with data as it enters Hive."<sup>2</sup>

Facebook's proposal is untenable for two reasons.

First, Exhibit B lists 11,051 Hive tables. Facebook proposes that 100 of those tables — fewer than 1% of them — be selected at random for limited production. While Facebook

---

<sup>2</sup> Facebook's April 11 submission states that "Exhibit B is a list of Hive tables that store user identifiers with data as it enters Hive," but then states that the tables "were identified using tools that probabilistically identified them as containing user identifiers."



Special Master Danie B. Garrie  
April 29, 2022  
Page 8

**KELLER ROHRBACK L.L.P.**  
**BLEICHMAR FONTI & AULD LLP**

proposes that Plaintiffs be permitted to select another 100 tables, that is not only (again) fewer than 1% of the tables it identified that may contain Named Plaintiff data, but Exhibit B does not contain sufficient information for Plaintiffs to make an informed determination about which tables to select.

Second, the information Facebook proposes to produce from those 200 tables is unduly limited. Facebook proposes that it “will produce a sample of data in warm storage” from the selected. That description is insufficiently specific. What does Facebook mean by “a sample of data”? What data is “in warm storage” (and, equally, what data is not)? Moreover, given the number of tables at issue and the hundreds of millions of users whose information may be contained therein, a sample of 200 tables is exceedingly unlikely to contain Named Plaintiff data. It is far too late in this already attenuated process and far too close to the end of fact discovery for this level of imprecision.

In addition, Plaintiffs are concerned about the extent of Facebook’s efforts to preserve Named Plaintiffs’ data stored in Hive. In its response to Plaintiffs’ inquiry about the time period it would produce, Facebook stated that the “[t]ime range depends on the tables selected.” Ex. 1. Plaintiffs need to understand what attempts Facebook made to retain data in Hive that can be associated with the Named Plaintiffs, as well as the standard retention policies for various tables.

Because Facebook’s proposal is untenable, Plaintiffs propose the following production protocol.

First, by May 6, 2022, Facebook should be ordered to describe the efforts it undertook to preserve data in Hive that can be associated with Named Plaintiffs and the timing of those efforts. On that date, Facebook should also provide the standard preservation practices applicable to each of the tables it identified in Exhibit B.

Second, by May 13, 2022, Facebook should produce to Plaintiffs and make available to the Special Master the schema and an excerpt comprising the top five rows of each of the Hive tables listed in Exhibit B. This information will help Plaintiffs decide the specific tables to search for Named Plaintiff data.

Third, by June 3, 2022, Plaintiffs will identify to Facebook and the Special Master the tables from which they seek Named Plaintiff data. If Facebook does not agree to produce from the tables identified by Plaintiffs, the parties will have one week to attempt to resolve the dispute, during which time they will be required to meet and confer at least once. If the parties have not reached agreement by June 10, the issue will be submitted, without additional briefing or explanation unless requested, to the Special Master.

Fourth, within one week of parties’ agreement on the tables to be searched or the Special Master resolves the issue, the Special Master will enter a rolling production schedule.

### 3. User data identified in Exhibit B of the April 11, 2022 submission and not in DYI

#### Facebook's Statement:

As shown in Exhibit C of Facebook's April 11, 2022 submission to the Special Master, only two types of user data in the contracts submitted to the Special Master for *in camera* review are not included in the DYI file: (1) privacy settings, and (2) set permissions (audience controls on a post). (1) has already been produced. Facebook will agree to produce (2) from TAO.

#### Plaintiffs' Response to User Data Identified in Exhibit C of April 11 Submission:

As an initial matter Plaintiffs believe Facebook intended to reference Exhibit C, not Exhibit B. The April 11 submission identified Exhibit C as "identif[ying] user data in the contracts submitted to Special Master Garrie for *in camera* review." Moreover, Plaintiffs' response assumes that Exhibit C accurately reflects all of the types of data in the contracts Facebook submitted to the Special Master for review.

Regarding privacy settings, attached to this response as Exhibits 3 and 4 are the privacy settings Facebook produced. Facebook should confirm that Exhibits 3 and 4 include all of the information Facebook has about each of the Named Plaintiff's privacy settings during the relevant time period. If there is any additional information related to the Named Plaintiffs' privacy settings, Facebook should produce it. Facebook should also describe any limitations on the privacy settings it is able to produce (e.g., it only retains privacy settings for a certain period of time, it does not retain privacy settings for Named Plaintiffs who are no longer on Facebook, etc.).

Regarding set permissions, Plaintiffs agree with Facebook's proposal so long as the production is not limited in any way. Facebook should also describe any limitations on the set permissions it is able to produce (e.g., it only retains set permissions for a certain period of time, it does not retain set permissions for Named Plaintiffs who are no longer on Facebook, etc.).

With respect to other Facebook systems, not including TAO and Hive, there are seven systems<sup>3</sup> that did not fall into one of the categories identified by the Special Master in Hearing Order Regarding Plaintiffs' Motion To Compel Production of Plaintiff Data, dated February 21, 2022: Ads Raw Storage, Manifold, orderdb (MySQL), Payments (MySQL), Tally, UDB, XDB, and ZippyDB. Producing data from any of these systems would be extremely burdensome and is not relevant or proportional to the needs of this case.

- *Ads Raw Storage and Manifold*: These systems are BLOB and binary storage and are not searchable by user identifiers.

---

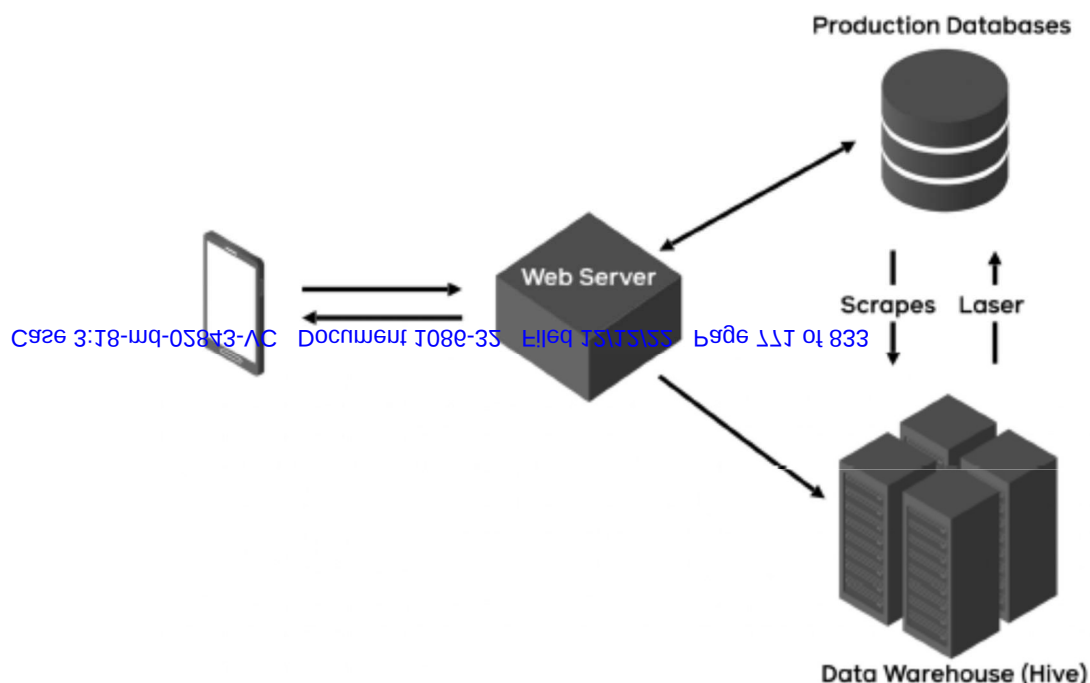
<sup>3</sup> Callisto should have been put in Category (6) because the data is in DYI. UDB should also have been put in Category (6) because it is duplicative of TAO. MySQL describes the general use of MySQL databases and should not be a separate entry. For example, orderdb (MySQL) and Payments (MySQL) are specific uses of MySQL.

- *Tally*: This system is a real-time aggregation counter and is not searchable by user identifiers.
- *orderdb (MySQL), Payments (MySQL), XDB, and Zippy DB*. These systems contain unstructured data and cannot be searched by user identifiers, except for tables using Ent or Node schemas for which there is a structured deletion plan. Searching other tables for named plaintiff data would require a full text search of all tables that would take many months to complete and generate a significant number of false positives requiring manual review of all data returned.

### Plaintiffs' Response:

Based on the expectation that Facebook's descriptions are accurate regarding what is contained in them and how that data is used and stored, Plaintiffs will not presently seek the data from the seven systems described in this section.

With respect to Dataswarm, the Special Master suggested that Facebook "query Tasks to identify Task Definitions which involve named plaintiff data" using identifiers, review those Task Definitions, and then search for named plaintiff data in the sources/destinations identified. We have investigated the Special Master's suggestion and did not include it in our proposed protocol because it is not feasible. Dataswarm is a collection of data processing operations coded in Python. Task Definitions are the Python code and Tasks are the running of that code. Searching Python code for user identifiers would be done through a syntax-only search which cannot understand the semantic meaning of the data being processed, i.e., whether a column name being processed involves a user identifier. As a result, additional code would have to be developed to understand whether/how identifiers are involved in a particular Task Definition.



Special Master Danie B. Garrie  
 April 29, 2022  
 Page 11

KELLER ROHRBACK L.L.P.  
 BLEICHMAR FONTI & AULD LLP

### Plaintiffs' Response:

Based on Facebook's description of Dataswarm, Plaintiffs will not presently seek production from that system.

### SCENARIOS

#### Special Master's Questions Following March 9, 2022 Hearing:

No later than April 4, 2022, Facebook is to submit documentation sufficient to describe the data collected both on and off platform or provided by Third Parties in the following scenarios and provide written responses to the questions below: "Exhibit A to Plaintiff's Questions re: Data Collection and Use indicates that Facebook used predictive algorithms to generate five political segments for Facebook users (Very Liberal, Liberal, Moderate, Conservative, and Very Conservative) based on demographic, psychographic, and behavioral signals from Facebook user data."

1. What are the inputs into these algorithms (i.e. what are the demographic, psychographic, and behavioral signals used to generate the political segments)? Are these inputs provided by users or derived by Facebook?

**Facebook's Answer:** As an initial matter, Very Liberal, Liberal, Moderate, Conservative, and Very Conservative are ad interests ("political ad interests"), which are distinct from political segmentation "based on demographic, psychographic, and behavioral signals from Facebook user data" ("political segments"). Political segments are political ad interests that are divided into political segments.

**Political ad targeting segments** were deprecated for new campaigns in January 2022 and for existing campaigns in March 2022. Based on a reasonable investigation to date, these were ad targeting segments derived from the following on-platform activity (user-provided data and observed data):

[REDACTED]

**Political segments** were deprecated for new campaigns in October 2017 and for existing campaigns in January 2018. Based on a reasonable investigation to date, these segments were created by

[REDACTED]

### Plaintiffs' Response:

Plaintiffs are concerned about the imprecision of Facebook's response. The use of phrases such as "related to *things like* device usage and travel preferences," "demographics *like*

Special Master Danie B. Garrie  
 April 29, 2022  
 Page 12

KELLER ROHRBACK L.L.P.  
 BLEICHMAR FONTI & AULD LLP

age, gender, and location,” and “[REDACTED]...” raises the possibility that Facebook has not responded to the Special Master’s inquiry by identifying all of the inputs. Further, Facebook’s response does not indicate how or what off-platform activity informed political ad targeting segments or political segments. No off-platform activities are mentioned as inputs into political ad targeting segments. And for political segments, “[REDACTED]” are identified without further description.

Facebook’s statement that political ad targeting segments were deprecated in March 2022, and political segments deprecated in January 2018 are non sequiturs. Plaintiffs’ proposed class period runs from 2007. Any responsive information from 2007 to the point of deprecation is relevant and should be provided.

2. How are the psychographic signals computed (e.g. how is the psychographic signal “High Dollar Religious Donor” determined)?

**Facebook’s Answer:** Based on a reasonable investigation to date, we believe “High Dollar Religious Donor” was an ad targeting option based on partner categories. We do not know how this partner category was created. Partner categories were deprecated in October 2018.

#### **Plaintiffs’ Response:**

Facebook does not answer the Special Master’s question. The only response it provides concerns an example identified by the Special Master for illustrative purposes, and even concerning that example, Facebook’s response is that it can’t respond because it doesn’t know.

Facebook’s statement that partner categories were deprecated in October 2018 is a non sequitur. Again, information from 2007 through deprecation is relevant to Plaintiffs’ claims.

3. Is information regarding identifiable ethnic affinities provided by users or derived by Facebook? How is ethnic affinity derived?

**Facebook’s Answer:** Ethnic affinities, also referred to as multicultural affinities, were ad interests created by Facebook based on a user’s on-platform activity indicating an interest in content relating to certain communities. Ethnic affinity ad interests were deprecated in August 2020.

#### **Plaintiffs’ Response:**

Facebook’s answer of how ethnic affinity derived is so general as to be almost meaningless. What “on-platform activity indicating an interest in content relate to certain communities” was considered? What does Facebook consider to be activity “indicating an interest in content relating to certain communities”? Which content? Which communities? How did Facebook use that information to derive ethnic affinities (or multicultural affinities)? What ethnic affinities were computed? Was any information other than the individual user’s on-platform activity indicating an interest in content related to certain communities considered?

Special Master Danie B. Garrie  
April 29, 2022  
Page 13

KELLER ROHRBACK L.L.P.  
BLEICHMAR FONTI & AULD LLP

Facebook's statement that ethnic affinity ad interests were deprecated in August 2020 is a non sequitur. Again, information from 2007 through deprecation is relevant to Plaintiffs' claims.

4. Where is political segmentation data for Facebook users stored?

**Facebook's Answer:** Political segments were deprecated in October 2017 for new campaigns. Existing campaigns using political segments ended on January 1, 2018. Based on a reasonable investigation, we believe that political segments were deleted in January 2018.

**Plaintiffs' Response:**

Plaintiffs are concerned by Facebook's decision to delete political segments data and ask that Facebook provide information about the circumstances that led to its deletion. Further, even if it is not retained, Plaintiffs ask Facebook to identify where political segments data was stored before it was deleted. That answer may help inform the parties' analysis of where relevant information associated with Named Plaintiffs is presently stored. Plaintiffs note that Facebook did not respond completely with respect to political segmentation data, as Facebook did not discuss political ad targeting segments in its response. Plaintiffs ask that Facebook supplement its response no later than May 6, 2022.

5. Is political segmentation determined for a Facebook user as part of a data process that runs on a regularly scheduled basis or evaluated in real time when an ad is served?

**Facebook's Answer:** Based on a reasonable investigation to date, we believe political segments were created manually by [REDACTED]

**Plaintiffs' Response:**

Facebook's response appears adequate with respect to political segments. Plaintiffs note, however, that Facebook did not respond completely with respect to political segmentation data, because it omitted political ad targeting segments from its response. Plaintiffs ask that Facebook supplement its response no later than May 6, 2022.

6. Is political segmentation associated with a Facebook user if possible (i.e. via UID, RID, SID, ASID, or another identifier that can be mapped to a user)? If so, explain how the political segmentation is associated with a Facebook user.

**Facebook's Answer:** Based on a reasonable investigation to date, we believe political segments were associated with users by UIDs.

**Plaintiffs' Response:**

Facebook's response appears adequate with respect to political segments. Plaintiffs note, however, that Facebook did not respond completely with respect to political segmentation data

because it omitted political ad targeting segments from its response. Plaintiffs ask that Facebook supplement its response no later than May 6, 2022.

7. Is an individual's assigned political segment part of the DYI file?

**Facebook's Answer:** Based on a reasonable investigation to date, we do not believe political segments were included in the DYI file.

**Plaintiffs' Response:**

This response conflicts with Facebook's April 11 submission. In its April 18 submission, Facebook defines political segments as political ad interests. In its April 11 submission, Facebook asserted that ad interests are "available in DYI." But in its April 18 submission, Facebook says that it does not believe political segments were included in the DYI file. Plaintiffs are concerned that there may be other ads interests that were also not included in the DYI file.

Plaintiffs also request that Facebook produce the assigned political segments and any other ad interests not included in the DYI file for each of the Named Plaintiffs.

Regards,



Derek W. Loeser  
[dloeser@kellerrohrback.com](mailto:dloeser@kellerrohrback.com)



Lesley E. Weaver  
[lweaver@bfalaw.com](mailto:lweaver@bfalaw.com)

# Exhibit R



**UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA  
SAN FRANCISCO DIVISION**

IN RE: FACEBOOK, INC. CONSUMER  
PRIVACY USER PROFILE LITIGATION,

This document relates to:

ALL ACTIONS

**MDL NO. 2843**

CASE NO. 3:18-MD-02843-VC-JSC

HON. VINCE CHHABRIA  
HON. JACQUELINE SCOTT CORLEY  
COURTROOM 4 – 17<sup>TH</sup> FLOOR  
SPECIAL MASTER, DANIEL GARRIE, ESQ.

**ORDER FOLLOWING MAY 17, 2022  
HEARING REGARDING PRODUCTION  
OF NAMED PLAINTIFF DATA**

**ORDER FOLLOWING MAY 17, 2022 HEARING REGARDING PRODUCTION OF NAMED  
PLAINTIFF DATA**

No later than June 2, 2022, Facebook is to submit responses to the following questions and provide the requested additional information.

- How did Facebook identify the Hive tables in Exhibit B to Facebook’s April 11, 2022 submission (“Exhibit B”) identified as possibly containing user identifiers? Facebook’s answer is to include technical details about the process used to identify the Hive tables.
- Can Facebook provide descriptions for the Hive tables identified in Exhibit B? If so, provide the descriptions.
- How long would it take to extract the column names and data types (schema) for the Hive tables identified in Exhibit B? Facebook’s answer is to include a description of the processes and steps it would need to perform to provide the requested data.
- Did Facebook’s search to identify the tables set out in Exhibit B include searching cold storage? If not, why not?
- In Exhibit B to Facebook’s April 18, 2022 letter submission (Declaration of Mengge Ji), Mengge Ji states “Tables are organized into more than 200 individual databases roughly based on product, team, or use case.” Expand on this statement, include descriptions of the 200+ databases, and clarify what is meant by “databases” in this context.
- One of Facebook’s responses in their April 11, 2022 letter submission at page 2 states “tables believed to contain data that that is either duplicative or transformed.” What does transformed mean in this context? Facebook’s answer is to include three examples of “transformation.”
- Facebook’s April 11, 2022 letter submission at page 2 refers to six categories of tables that were not included in the list in Exhibit B: (1) temporary, test, and error tables; (2) tables used to evaluate system performance; (3) tables that contain no data; (4) tables related to other Meta products (e.g., Oculus, Instagram, and WhatsApp); (5) tables used for infrastructure, operations,


1 and product testing; and (6) tables believed to contain data that is either duplicative or  
2 transformed. Provide a description of each of these six categories of tables.

3 No later than May 30, 2022, Facebook is to provide the revised proposal they discussed at the  
4 hearing.

5 No later than June 6, 2022, the parties are to submit briefs, not to exceed five (5) pages not  
6 including exhibits, on the issues of Hive table selection, production of data related to apps installed by  
7 friends of Named Plaintiffs, and whether Facebook should be compelled to search cold storage for  
8 Named Plaintiff data.  
9

10  
11 IT IS SO ORDERED.

12  
13 Tuesday, May 24, 2022

14   
Daniel Garrie  
Discovery Special Master

# Exhibit S

GIBSON DUNN

Gibson, Dunn & Crutcher LLP  
555 Mission Street  
San Francisco, CA 94105-0921  
Tel 415.393.8200  
www.gibsondunn.com

Rosemarie T. Ring  
Direct: +1 415.393.8247  
Fax: +1 415.801.7358  
RRing@gibsondunn.com

HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

May 30, 2022

VIA JAMS ACCESS

Special Master Daniel B. Garrie  
DGarrie@jamsadr.com

Re: *In re Facebook Consumer Privacy User Profile Litigation*, JAMS Ref No.  
1200058674

Dear Special Master Garrie,

Pursuant to the Special Master's Order Following May 17, 2022 Hearing Regarding Production of Named Plaintiff Data, Facebook submits this updated proposal re: named plaintiff data, reflecting additional terms Facebook has agreed to in response to requests from Plaintiffs in response to Facebook's initial proposal, to assist the Special Master in determining "what, if any" additional named plaintiff data should be produced, as directed in Judge Corley's order dated January 12, 2022 (Dkt. No. 807).

**Facebook's Proposal For Producing Additional Named Plaintiff Data**

As Judge Corley explained in her most recent order on named plaintiff data, issued on January 12, 2022, the Special Master is working with the parties to determine "what, if any" additional named plaintiff data should be produced consistent with Rule 26. Facebook continues to believe that the named plaintiff data that has already been produced, which includes data in all three categories of "discoverable user data" identified in Judge Corley's Discovery Order No. 9, satisfies its obligations under Rule 26. The additional data Plaintiffs seek is neither relevant to their claims nor proportional to the needs of this case because it was not shared or made accessible to third parties.

That said, in an effort to resolve this issue, and consistent with Judge Corley's Discovery Order Nos. 11 and 12 allowing Plaintiffs to test Facebook's position on whether there is additional named plaintiff data that was shared or made accessible to third parties, Facebook proposes the protocol set forth below for producing additional named plaintiff data.

## GIBSON DUNN

Special Master Daniel B. Garrie  
May 30, 2022  
Page 2

### HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

As explained in Facebook’s March 7, 2022 submission to the Special Master, at a high level, user data is stored in production systems and the data warehouse. For production systems we look to TAO, and for the data warehouse we look to Hive. Accordingly, we propose to produce additional data from these systems as explained below.

**TAO.** As also explained in the March 7, 2022 submission, TAO is a distributed data store for the social graph. There is an object and associations to that object for all users. *See* Ex. A, Internal Facebook Wiki Regarding TAO Core Concepts. For each named plaintiff, Facebook will produce the user objects and associations to those objects.

**May 30, 2022 Update:** Facebook will also provide TAO schema for the above data.

**Hive.** As explained in prior submissions to the Special Master, Hive is Facebook’s data warehouse. *See* Ex. B, Declaration of Mengée Ji In Support of Facebook Inc.’s Motion For Reconsideration of Special Master’s Order Regarding Named Plaintiff Data, ¶ 11. Hive is not structured to support centralized searches for individual user data. *Id.* ¶ 16. Instead, it is organized roughly based on product, team, or use case. *Id.* ¶ 17. Searching for specific types of data typically requires seeking the advice of product teams or other subject matter experts to identify potentially relevant tables. *Id.* ¶ 18.

Once potentially relevant tables are identified, they must be individually assessed to determine whether they contain user data and then searched for data related to the named plaintiffs. *Id.* These searches are subject to many technical and policy limitations which make searching for individual user data unduly burdensome, including that (1) many Hive tables do not associate data with individual users (*id.* ¶ 20); (2) many Hive tables are stored in “partitions” or segments, and many of those tables are partitioned by day, reflecting the time when the data was logged or processed, such that each day of data is its own segment and must be independently searched (*id.* ¶ 21); (3) data in many Hive tables is machine readable (*id.* ¶ 22); and (4) data in many Hive tables are in cold storage and would have to be restored in order to be searched and analyzed, *see* Ex. C, Declaration of Mengée Ji In Support of Facebook Inc.’s Motion For a Protective Order Against Production of API Call Logs, ¶¶ 17-18; Special Master’s Nov. 8, 2021 Order Re: Facebook’s Motion For Protective Order Against Production of API Call Logs, ¶ 15 (“Special Master Garrie finds that the data in the Mobile Table and Web Table is not reasonably accessible because it is not readily usable in its ‘cold storage’ state and must be restored to ‘warm storage’ in order to be searched and analyzed (i.e. usable).”).

For these reasons, among others, as reported in its April 11, 2022 submission, Facebook estimates that it would take approximately 62 billion CPU seconds or around

# GIBSON DUNN

Special Master Daniel B. Garrie  
May 30, 2022  
Page 3

## HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

720,000 CPU days of computational cost to search across the entire Hive data warehouse and extract all data about a single user.

Once the searches have been conducted, and data is returned, the data must be manually reviewed to, among other things, confirm that the data is associated with a particular user and does not include personal data of other users (e.g., user blocking another user) or trade secrets, and does not create system integrity or security concerns.

In light of the above technical limitations and burdens, Facebook proposes to search for and produce the following named plaintiff data from Hive to the extent it exists: (1) specific types of data requested or referenced by Plaintiffs in challenging Facebook’s production of named plaintiff data; (2) a sample of data from Hive tables with user identifiers included in Exhibit B to the April 11, 2022 submission; and (3) the only data identified in Exhibit C to the April 11, 2022 submission that has not already been produced to Plaintiffs in the DYI file or otherwise: “set permissions (audience controls on a post).”<sup>1</sup>

### 1. Specific Types of Data

- *Off-platform activity*: DYI includes off-platform activity. Facebook will search for and produce any underlying raw log-level data for off-platform activity provided to Facebook by a third party associated with the named plaintiffs.
- *Ad interests*: DYI includes ad interests. Facebook will search for and produce any underlying raw log-level data associated with the named plaintiffs.
- *Ad click data*: DYI includes ad click data. Facebook will search for and produce any underlying raw log-level ad clicks data associated with the named plaintiffs.
- *Ad impressions data*: Facebook will search for and produce ad impressions data associated with the named plaintiffs.
- *Custom audience data*: DYI identifies third parties who have created custom audiences associated with a user. Facebook will search for and produce any more granular information about custom audiences associated with the named plaintiffs, including custom audience type, whether it was used to deliver ads, and when.

---

<sup>1</sup> The contract referencing this data allowed a user to “set permissions (audience controls on a post)” from their phone. Facebook does not use the term “set permissions” to refer to this data.

## GIBSON DUNN

Special Master Daniel B. Garrie  
May 30, 2022  
Page 4

HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

**May 30, 2022 Updates:** Facebook responds as follows to plaintiffs' additional requests:

- *Scope of production.* Facebook confirms it will produce the above data regardless of whether it appears in DYI. Facebook also confirms its productions include data “that can be associated” with each named plaintiff and notes there is no distinction in Hive between data that is associated with a user and data that can be associated with a user. Hive data that is or can be associated with a user has a user identifier.
- *Tables and schema.* Facebook will provide the names of the tables from which the five categories of Hive data described above will be produced and how it identified the tables, as well as the schema for the above data.
- *Data described by Aldo King.* Facebook has already agreed to search for and produce ad interests in Hive, regardless of whether it appears in DYI.
- *Installed apps.* Apps installed by a user through Facebook Login (either on or off of Facebook) are listed in DYI, and information about apps installed by users without Facebook Login appears in the Off-Facebook Activity section of DYI.

### **2. Sampling from Hive tables identified in Exhibit B of April 11, 2022 submission**

Given that data in Hive tables are not shared or made accessible to third parties, the significant burden of searching for and producing individual user data from Hive as explained in prior submissions, and that Facebook has already produced nearly a million pages of the named plaintiff data from the DYI system, Hive data is not relevant or proportional to the needs of this case. Facebook nevertheless understands that, as Judge Corley found in Discovery Order Nos. 11 and 12, Plaintiffs are allowed to test Facebook's positions on “sharing and accessibility” by obtaining discovery into what data exists and how it is used. To that end, and consistent with the requirements of relevance and proportionality under Rule 26, Facebook proposes a sampling process.

**May 30, 2022 Update:** In addition to the above specific categories of data, Facebook will produce named plaintiff data in warm storage for a statistically significant sample of Hive tables identified in Exhibit B of the April 11, 2022 submission. Specifically, Facebook will produce named plaintiff data from 500 tables in Exhibit B to be jointly selected by Facebook and Plaintiffs. Facebook will randomly select 250 tables and Plaintiffs will select 250 tables.



GIBSON DUNN

Special Master Daniel B. Garrie  
May 30, 2022  
Page 5

HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

**3. User data identified in Exhibit C of the April 11, 2022 submission and not in DYI**

As shown in Exhibit C of Facebook’s April 11, 2022 submission to the Special Master, only two types of user data in the contracts submitted to the Special Master for in camera review are not included in DYI: (1) privacy settings, and (2) “set permissions (audience controls on a post.” (1) has already been produced. Facebook will agree to produce (2), including any audience controls on individual pieces of content, from TAO.

**May 30, 2022 Update:** Facebook produced account-level privacy settings from the named plaintiffs’ accounts in 2020. Facebook will produce updated privacy settings for each named plaintiff.

Sincerely,

A handwritten signature in cursive script, reading "Rosemarie Ring".

Rosemarie T. Ring

# Exhibit T

GIBSON DUNN

Gibson, Dunn & Crutcher LLP  
555 Mission Street  
San Francisco, CA 94105-0921  
Tel 415.393.8200  
www.gibsondunn.com

Rosemarie T. Ring  
Direct: +1 415.393.8247  
Fax: +1 415.801.7358  
RRing@gibsondunn.com

HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

June 2, 2022

VIA JAMS ACCESS

Special Master Daniel B. Garrie  
DGarrie@jamsadr.com

Re: *In re Facebook Consumer Privacy User Profile Litigation*, JAMS Ref No.  
1200058674

Dear Special Master Garrie,

On behalf of Facebook, below we submit answers to the questions in the Special Master's Order Following May 17, 2022 Hearing Regarding Production of Named Plaintiff Data, dated May 24, 2022.

- 1. How did Facebook identify the Hive tables in Exhibit B to Facebook's April 11, 2022 submission ("Exhibit B") identified as possibly containing user identifiers? Facebook's answer is to include technical details about the process used to identify the Hive tables.**

Facebook created the list of tables in Exhibit B to Facebook's April 11, 2022 submission using a system that performs data classification in Facebook's data warehouse. The system takes a sample of the data in a table and classifies the type of data. Tables are scanned periodically to refresh their prediction.

The classification system is able to classify tables with UID, RID and SID. The data classification system first identifies if a column value is a number. If so, it determines the type of identifier, if any.

The data classification analysis used for this exercise included tables in cold storage. The classification system can identify whether a table in cold storage contains a UID, RID and SID. The classification system does not search within tables in cold storage for a particular individual user's UID, RID and SID nor does it restore data into the live data warehouse.

## GIBSON DUNN

Special Master Daniel B. Garrie  
June 2, 2022  
Page 2

HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

- 2. Can Facebook provide descriptions for the Hive tables identified in Exhibit B? If so, provide the descriptions.**

Facebook will provide plaintiffs the column names for the Hive tables identified in Exhibit B to Facebook's April 11, 2022 submission by June 9, 2022. This data is also the schema for a Hive table. Extracting this information requires running a query for the columns in all of the tables listed in Exhibit B.

- 3. How long would it take to extract the column names and data types (schema) for the Hive tables identified in Exhibit B? Facebook's answer is to include a description of the processes and steps it would need to perform to provide the requested data.**

See response to Question 2.

- 4. Did Facebook's search to identify the tables set out in Exhibit B include searching cold storage? If not, why not?**

As noted, the list of tables in Exhibit B to Facebook's April 11, 2022 submission includes tables in cold storage.

- 5. In Exhibit B to Facebook's April 18, 2022 letter submission (Declaration of Mengge Ji), Mengge Ji states "Tables are organized into more than 200 individual databases roughly based on product, team, or use case." Expand on this statement, include descriptions of the 200+ databases, and clarify what is meant by "databases" in this context.**

Facebook organizes Hive tables into namespaces, which are essentially separate databases, primarily for capacity-budgeting reasons. These namespaces are attached as Exhibit A.

- 6. One of Facebook's responses in their April 11, 2022 letter submission at page 2 states "tables believed to contain data that is either duplicative or transformed." What does transformed mean in this context? Facebook's answer is to include three examples of "transformation."**

As the Special Master observed, the term "transformed data" is standard SQL terminology. Data transformation is the process of converting data from one format to another (such as by changing, removing, or combining data) and is part of a typical data

## GIBSON DUNN

Special Master Daniel B. Garrie  
June 2, 2022  
Page 3

### HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

management process, including data warehousing. Three examples of transformation are below:

- (1) A table is transformed when data within it is de-identified and/or pseudonymized.

[REDACTED]

In this scenario, the transformed version of the table was not included on the list of tables in Exhibit B to Facebook's April 11, 2022 submission to not overstate the number of tables or list duplicative tables.

[REDACTED]

- (2) A table is transformed when it combines data from multiple tables. One table may store certain demographic information about users. Another table may store certain user actions. A transformed table may combine the data from those two tables to reflect both user demographic information and user actions.

- (3)

[REDACTED]

7. **Facebook's April 11, 2022 letter submission at page 2 refers to six categories of tables that were not included in the list in Exhibit B: (1) temporary, test, and error tables; (2) tables used to evaluate system performance; (3) tables that contain no data; (4) tables related to other Meta products (e.g., Oculus, Instagram, and WhatsApp); (5) tables used for infrastructure, operations, and product testing; and (6) tables believed to contain data that is either duplicative or transformed. Provide a description of each of these six categories of tables.**

- (1) Temporary, test, and error tables.

Facebook excluded temporary tables that are created to perform queries without needing to recompute data. Facebook similarly excluded tables an engineer used to test changes to the table, as these are duplicative of the original table. Facebook excluded tables that were created because of rare errors in the processing of data to the original table.

GIBSON DUNN

Special Master Daniel B. Garrie  
June 2, 2022  
Page 4

HIGHLY CONFIDENTIAL—ATTORNEYS EYES ONLY

- (2) Tables used to evaluate system performance and (5) tables used for infrastructure, operations, and product testing.

Facebook excluded tables that are used to evaluate system performance and/or infrastructure, operations, and product testing. Infrastructure, system performance, and operations tables are used for the operational health of Facebook's systems, such as error logging, performance logging of code, and network performance. Product testing tables are used internally by Facebook product teams in order to test and build new features and products.

- (3) Tables that contain no data.

Facebook excluded any table that does not include data. As an example, a table an engineer created but to which no data has been added is a table containing no data.

- (4) Tables related to other Meta products (e.g., Oculus, Instagram, and WhatsApp).

Facebook excluded tables from namespaces that relate only to Meta products other than Facebook, such as the instagram, whatsapp, and oculus namespaces.

- (6) Tables believed to contain data that is either duplicative or transformed.

Facebook excluded tables that contain duplicative or transformed data. Duplicative data is data that is the same as data in a different table. Transformed data is described above in response to Question 6.

Sincerely,

A handwritten signature in cursive script that reads "Rosemarie Ring".

Rosemarie Ring

# Exhibit A

Highly Confidential –  
Attorneys' Eyes Only

[REDACTED]



[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

# Exhibit U

GIBSON, DUNN & CRUTCHER LLP  
Orin Snyder (*pro hac vice*)  
osnyder@gibsondunn.com  
200 Park Avenue  
New York, NY 10166-0193  
Telephone: 212.351.4000  
Facsimile: 212.351.4035

Kristin A. Linsley (SBN 154148)  
klinsley@gibsondunn.com  
Rosemarie T. Ring (SBN 220769)  
rring@gibsondunn.com  
Martie Kutscher (SBN 302650)  
mkutscherclark@gibsondunn.com  
555 Mission Street, Suite 3000  
San Francisco, CA 94105-0921  
Telephone: 415.393.8200  
Facsimile: 415.393.8306

*Attorneys for Defendant Facebook, Inc.,*

GIBSON, DUNN & CRUTCHER LLP  
Deborah Stein (SBN 224570)  
dstein@gibsondunn.com  
Heather L. Richardson (SBN 246517)  
hrichardson@gibsondunn.com  
333 South Grand Avenue  
Los Angeles, CA 90071-3197  
Telephone: 213.229.7000  
Facsimile: 213.229.7520

Joshua S. Lipshutz (SBN 242557)  
jlipshutz@gibsondunn.com  
1050 Connecticut Avenue, N.W.  
Washington, DC 20036-5306  
Telephone: 202.955.8500  
Facsimile: 202.467.0539

**UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA  
SAN FRANCISCO DIVISION**

IN RE: FACEBOOK, INC. CONSUMER  
PRIVACY USER PROFILE LITIGATION,

This document relates to:

ALL ACTIONS

CASE NO. 3:18-MD-02843-VC

**FACEBOOK, INC.'S BRIEF  
REGARDING THE PARTIES' NAMED  
PLAINTIFF DATA PROPOSALS**

As Judge Corley explained in her most recent order on this issue, the purpose of these proceedings is to determine “what, if any, data from [systems other than DYI] should be produced consistent with [FRCP 26(b)].” Dkt. 807 at 4. Facebook continues to believe that its productions to date, which include all three categories of “discoverable user data” identified by Judge Corley in Discovery Order No. 9 (Dkt. 557), meet this standard, but submitted a proposal agreeing to search for and produce additional data. To further facilitate resolution of the named plaintiff data issue, Facebook agreed to most of Plaintiffs’ demands in their response to the proposal and offered reasonable compromises. During the May 17, 2022 hearing to discuss Facebook’s proposal and Plaintiffs’ response, the Special Master identified three remaining issues: (1) Hive data, (2) apps installed by friends of the named plaintiffs, and (3) cold storage. As explained below, Facebook will agree to provide (2), and the Special Master should accept Facebook’s proposed compromise on (1) and defer (3) as premature.

#### **1. Hive Tables**

Facebook proposes to search for and produce named plaintiff data from a sample of the 11,051 Hive tables identified in Exhibit B to its April 11, 2022 submission. Initially, Facebook proposed a sample of 200 tables, with Plaintiffs selecting 100 tables and Facebook randomly selecting 100 tables. In response, Plaintiffs demanded that Facebook produce all data (regardless of whether it is named plaintiff data) in the first five rows of all 11,051 Hive tables. Based on this data, Plaintiffs would identify what additional data they want, the parties would have one week to resolve any disputes, and, if no agreement were reached, the issues would go to the Special Master without briefing. During the May 17, 2022 hearing, Facebook offered a compromise which more than doubled its proposed sample, from 200 to 500 tables, representing nearly 5% of the 11,051 tables, with Plaintiffs selecting 250 tables and Facebook randomly selecting 250 tables. To the extent Plaintiffs’ proposal is seeking information about the contents of these Hive tables, Facebook agreed in its June 2, 2022 submission to provide Plaintiffs with the schema (column names) for all 11,051 Hive tables from which the sample would be selected by June 9, 2022.

It is worth noting, again, that Hive data is not accessible to third parties, and therefore is irrelevant to Plaintiffs’ claims in this case which are about data sharing. But in the interest of



1 compromise and to assist the Special Master in resolving this issue, Facebook has agreed to search  
 2 for and produce named plaintiff data from what the Special Master has generally treated as  
 3 representing a statistically significant sample (5%) of Hive tables.

4 Plaintiffs' proposal in response is unworkable. As an initial matter, it does not even seek  
 5 named plaintiff data. Plaintiffs want all data in the "first five rows" of all 11,051 Hive tables.  
 6 Producing this enormous volume of data, most of which is unlikely to include any named plaintiff  
 7 data, is not relevant or proportional to the needs of this case, and is unduly burdensome. The only  
 8 reason Plaintiffs offered at the May 17, 2022 hearing for seeking this data was to help Plaintiffs  
 9 understand the type of data contained in the 11,051 Hive tables. Facebook's agreement to provide  
 10 Plaintiffs with the schema (column names) for all 11,051 Hive tables from which the sample would  
 11 be selected fully addresses that supposed need.

12 Plaintiffs' proposal would take months to complete. First, Facebook would have to extract,  
 13 review, and produce all data in the top five rows of all 11,051 tables. Then, three weeks later,  
 14 Plaintiffs would identify an unspecified number of tables from which Facebook would search to  
 15 determine whether they contain named plaintiff data, produce any such named plaintiff data, meet  
 16 and confer with Plaintiffs over any disputes, and then bring any disputes to the Special Master for  
 17 resolution. Depending on the number of tables Plaintiffs choose, it could take months for disputes to  
 18 be resolved and for Facebook to begin producing data. Once the data extraction process begins, it  
 19 could take many more months to actually collect the data (depending on the scope of the requested  
 20 tables). This process is cumbersome, time consuming, and unduly burdensome given that it is also  
 21 extremely unlikely to provide Plaintiffs with any relevant information, much less information that is  
 22 probative of their claims.

23 Facebook's proposal, by contrast, provides Plaintiffs with named plaintiff data from a  
 24 statistically significant sample of Hive tables (500 out of 11,051 tables), half of which would be  
 25 selected by Plaintiff using the schema (column names) for all 11,051 tables.

## 26 **2. Apps Installed by *Friends* of the Named Plaintiffs**

27 At the May 17, 2022 hearing, Plaintiffs asked Facebook to produce a list of apps installed by  
 28 friends of the named plaintiffs. In Discovery Order No. 9, Judge Corley identified three categories of

“discoverable user data”: (1) data collected from a user’s on-platform activity, (2) data obtained from third parties regarding a user’s off-platform activities, and (3) data inferred from a user’s on or off-platform activity. Apps used by the friends of the named plaintiffs do not fall into any of these categories of data, and therefore are beyond the scope of these proceedings. That said, in the interest of facilitating resolution of this issue, Facebook will provide Plaintiffs data regarding interactions that friends of the named plaintiffs had with businesses/apps using Facebook Login, without identifying which friend interacted with each business/app.

### 3. Data Cold Storage

As Facebook explained in its June 2, 2022 submission, the 11,051 tables identified in Exhibit B to its April 11 submission include tables with data in cold storage (including tables with data only in cold storage). A table itself is not “in cold storage” or “in warm storage.” Data within a table is in cold or warm storage, and a single table can contain some data in warm storage and some data in cold storage. There is no need to restore data in cold storage to investigate whether tables contain user data. As explained in Facebook’s June 2, 2022 submission, the data classification system used to identify the 11,051 tables in Exhibit B to Facebook’s April 11 submission included tables with data in cold storage.

On November 18, 2021, the Special Master ruled that data in cold storage did not need to be produced because it had to be restored to warm storage in order to be searched and therefore was not “reasonably accessible.” *See* Special Master’s Nov. 8, 2021 Order Re: Facebook’s Motion For Protective Order Against Production of API Call Logs ¶ 15. The same is true here. That said, Facebook is willing to consider restoring data from cold storage, but identifying what, if any, data should be restored is premature at this stage. Facebook proposes that the parties first identify which Hive tables will be produced. Once the Hive tables have been identified, Facebook will evaluate which tables, if any, include data in cold storage, and make a proposal regarding what, if any, data it can reasonably restore, search, and produce.

1 Dated: June 6, 2022

**GIBSON, DUNN & CRUTCHER, LLP**

2  
3 By: /s/ Rosemarie T. Ring  
Orin Snyder (*pro hac vice*)  
osnyder@gibsondunn.com  
4 200 Park Avenue  
New York, NY 10166-0193  
5 Telephone: 212.351.4000  
Facsimile: 212.351.4035

6  
7 Deborah Stein (SBN 224570)  
dstein@gibsondunn.com  
8 Heather L. Richardson (SBN 246517)  
hrichardson@gibsondunn.com  
9 333 South Grand Avenue  
Los Angeles, CA 90071-3197  
10 Telephone: 213.229.7000  
Facsimile: 213.229.7520

11 Joshua S. Lipshutz (SBN 242557)  
jlipshutz@gibsondunn.com  
12 1050 Connecticut Avenue, N.W.  
Washington, DC 20036-5306  
13 Telephone: 202.955.8500  
Facsimile: 202.467.0539

14  
15 Kristin A. Linsley (SBN 154148)  
klinsley@gibsondunn.com  
16 Rosemarie T. Ring (SBN 220769)  
rring@gibsondunn.com  
17 Martie Kutscher (SBN 302650)  
mkutscherclark@gibsondunn.com  
18 555 Mission Street, Suite 3000  
San Francisco, CA 94105-0921  
19 Telephone: 415.393.8200  
Facsimile: 415.393.8306

20 Russell H. Falconer (*pro hac vice*)  
rfalconer@gibsondunn.com  
21 2100 McKinney Ave., Suite 1100  
Dallas, TX 75201  
22 Telephone: 214.698.3170  
Facsimile: 214.571.2958

23  
24 *Attorneys for Defendant Facebook, Inc.*

# Exhibit V

Derek W. Loeser (admitted *pro hac vice*)  
KELLER ROHRBACK L.L.P.  
1201 Third Avenue, Suite 3200  
Seattle, WA 98101  
Tel.: (206) 623-1900  
Fax: (206) 623-3384  
dloeser@kellerrohrback.com

Lesley E. Weaver (SBN 191305)  
BLEICHMAR FONTI & AULD LLP  
555 12th Street, Suite 1600  
Oakland, CA 94607  
Tel.: (415) 445-4003  
Fax: (415) 445-4020  
lweaver@bfalaw.com

*Plaintiffs' Co-Lead Counsel*

*[Additional counsel listed on signature page]*

**UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA  
SAN FRANCISCO DIVISION**

IN RE: FACEBOOK, INC. CONSUMER  
PRIVACY USER PROFILE LITIGATION

MDL No. 2843  
Case No. 18-md-02843-VC

This document relates to:

ALL ACTIONS

**PLAINTIFFS' RESPONSE TO SPECIAL  
MASTER'S ORDER OF MAY 24 RE:  
NAMED PLAINTIFF DATA**

Judge: Hon. Vince Chhabra  
Special Master Daniel Garrie  
Courtroom: 4, 17th Floor

JAMS Ref. No.: 1200058674

## I. BRIEF OVERVIEW

The Special Master’s May 24, 2022 Order requested briefing on three topics: the issues of Hive table selection; whether Facebook should be compelled to search cold storage; and production of data related to apps installed by friends of Named Plaintiffs.

Since the Special Master first requested this brief, new evidence has come to light in two 30(b)(6) depositions related to those questions: a deposition on evidence preservation (“preservation deposition”) and a deposition on Facebook’s ability to associate and delete users’ data through identifiers (“Topic 4 deposition”). The new evidence shows that: (1) Facebook selected 137 Hive tables and put them in “cold storage” precisely because they were relevant to this litigation;<sup>1</sup> (2) Facebook is capable of searching offline Hive tables using HQL and the Scuba tool; (3) the DYI file is not the most complete or usable compilation of user data; and (4) Facebook has withheld from production at least 52 snapshots of Named Plaintiff data using a never-before revealed tool more commonly used to collect user data called Switchboard, which preserves privacy settings and relationships between users and friends. These revelations cannot be reconciled with Facebook’s representations throughout the course of this litigation, which have significantly delayed discovery of Named Plaintiffs’ data and this litigation more generally.

### A. Facebook Hid Evidence of 137 Hive Tables Relevant to This Litigation Preserved in Cold Storage

The preservation deponent, Michael Duffey, testified that Facebook put 137 Hive tables in “cold storage” as a result of the litigation hold issued in this case. Duffey 30(b)(6) Dep. Rough Tr. at 60:21-25, 64:2-3.<sup>2</sup> This, presumably, is “offline.” The very purpose of putting this data into “cold storage” was to preserve them because they are relevant to this case. *Id.* at 66:18-67:1. For example, Mr. Duffey repeatedly pointed to these tables as a potential source of information relating to the data Facebook shared with third parties about users. *See, e.g., id.* at 148:1-3, 148:14-17,

---

<sup>1</sup> This raises a question as to whether “cold storage” truly makes them inaccessible. Why would Facebook make data *less* accessible if it were preserving it for litigation?

<sup>2</sup> Mr. Duffey was unable to answer many specific questions about these tables, including what they are; why they were selected; when they were put into cold storage; and whether they have ever been searched. Duffey 30(b)(6) Rough Tr. at 60:11, 60:21-25, 61:22, 65:3-4. Plaintiffs request production of a 30(b)(6) deponent, Jennifer Allen, on these topics.

153:14-17. Mr. Duffey testified that Facebook’s E-Discovery team still has access to these tables while they are in cold storage, that Facebook can produce the structured data from cold storage, and in fact, it has done so on previous occasions. *Id.* at 174:9, 175:7-13, 176:11-13. Facebook told the Special Master that this could not be done. *See* Declaration of Mengge Ji in Supp. of Facebook Inc.’s Mot. for a Protective Order Against Production of API Call Logs (“Ji Decl.”), ¶ 18 (“Tables in cold storage cannot be reviewed or analyzed while they remain in cold storage.”).

The Topic 4 deponent, Mike Clark, who previously appeared in a Special Master hearing, testified that offline Hive tables can be queried using HQL and the Scuba tool. Clark 30(b)(6) Dep. Tr., 86:20-89:23. Clark further testified that Hive tables can be searched either by UID or RID. *Id.* at 99:9-18. The table that contains the mapping of FBID to RIDs, `dim_fbuid_to_rid`, is available and queryable via TAO. *Id.* at 101:16-102:16. If Facebook knows the table and its location, it can query the specific table, including for device IDs. *Id.* at 114:10-14. (Ex. 334 to the Clark 30(b)(6) Dep. is a multipage document from 2010 that “contains all of the recently used Hive tables as defined by any table which has at least one user in the past 30 days as of 5/12/2010.”). Plainly, Facebook has the ability to produce summary documents of Hive tables with their fields (which includes UIDs, among other identifiers). But Facebook has repeatedly refused to conduct any searches of Hive, using as an impediment that it would be impossible to search all of Hive and Facebook cannot identify which tables to search. *See* Exhibit A.

**B. Facebook Improperly Withheld Snapshots of User Data Preserved with the Switchboard Tool, Among Other Data.**

For years, Facebook has insisted that the DYI files it has produced contain the “most complete compilation of data associated with the Named Plaintiffs’ accounts.” *See* Exhibit A. Facebook fiercely resisted discovery into whether this was true. *See, e.g.*, Facebook’s November 15, 2021 Response to Objection Regarding Named Plaintiffs’ Data Briefing (falsely stating that the DYI file was comprehensive and attacking Plaintiffs for probing that representation).

Mr. Duffey’s testimony, however, indicates that the DYI files are *not* complete. He explained that “Switchboard” snapshots contain information not captured by DYI, including information about users’ involvement in Pages, Groups, and advertising accounts. Duffey 30(b)(6)

Dep. Rough Tr. at 31:11-13; *see also* Exhibit 386 at -00003258. Critically, he testified that *they contain information about privacy settings*, which Facebook only recently admitted is not in the DYI file, as well as information about the “relationships between a user”—i.e., between one user and another—such as information about whom the user is interacting with. Duffey 30(b)(6) Dep. Rough Tr. at 45:7-10, 56:5-12. Mr. Duffey also testified that the DYI file is not the file that Facebook regularly uses to produce data about users in response to subpoenas. Rather, it is the Switchboard tool. *Id.* at 32:21-33:6. Switchboard snapshots are produced in a more “usable way than DYI” and is similar to a “PDF record.” *Id.* at 33:2-9. Facebook recorded Switchboard snapshots of the Named Plaintiffs on March 9, 2020 and in the fall of 2021, long before the Special Master held the first hearing on the Named Plaintiffs’ data. *Id.* at 50:16-20; *see also* Ex. 386 at -00003258.

Plaintiffs have long requested that more usable versions of DYI files be produced and sought more information about unproduced data. Dkt. No. 526 at 7 (noting the DYI file is incomplete and that format of production “obscures” information about the Named Plaintiffs’ privacy settings). Mr. Duffey’s testimony makes clear that Facebook has (1) falsely testified that the DYI file was the most complete collection of user data; (2) failed to inform the Court of the existence of the tool most commonly used to produce user data; (3) moved 137 Hive tables to “cold storage” precisely because they are relevant to this litigation, without revealing their existence; and (4) proposed an arduous and unrealistic process for Plaintiffs and the Special Master to guess at which Hive tables might possibly contain relevant data, including a “random” selection of 500 tables, while failing to disclose that Facebook had already identified 137 tables as relevant and preserved them offline.

## II. PLAINTIFFS’ INITIAL PROPOSAL

### A. Facebook Must Immediately Search the 137 Hive Tables with All Identifiers, Including Device IDs, Associated with the Named Plaintiffs.

Facebook should have voluntarily disclosed the existence of the 137 Hive tables during ESI discussions in 2019. At the least, it should have disclosed the tables to the Special Master.



Consider the direct questions posed by the Special Master over the past six months and compare Facebook's evasive responses. The Special Master ordered Facebook to identify "a list of data sources that may contain information related to the Named Plaintiffs," and, more specifically: "[i]dentify a list of Hive tables containing columns that store a UID, RID, SID, ASID, or other means of identifying a Facebook user. The list is to include tables that were active during the relevant time period." *See* Exhibit A. Facebook artfully answered other questions and did not identify the 137 tables. *Id.* Even Facebook's May 30 data production proposal did not disclose the existence of those tables, although it indicated with no specificity that cold storage tables are included in Exhibit B. The May 30 submission proposed only producing data in "warm storage for a statistically significant sample of Hive tables."<sup>3</sup> So, even though Facebook knew that it had already determined that relevant Hive tables exist, and had preserved them, it repeatedly proposed that the parties and the Special Master play a guessing game about whether Hive tables contained relevant Named Plaintiff data. All the while, it knew that the potentially most relevant tables would be excluded from the tables to be searched. Duffey 30(b)(6) Dep. Rough Tr. at 74:18-23, 60:21-25, 153:14-17. Months have passed as this inquiry has dragged on.

Furthermore, Mr. Clark's testimony makes perfectly clear that "offline" tables can be searched. Mr. Duffey also appeared to agree that such tables can be searched, using tools such as Scuba and iData. Duffey 30(b)(6) Dep. Rough Tr. at 61:5-8, 176:11-13; Ex. 386 at -00003259 ("iData: Allows searching for Hive tables and indicates if table is on hold."). Yet Facebook has opposed searching these records and stated that while in cold storage, data cannot be reviewed or analyzed. *See* Ji Decl., ¶ 18.<sup>4</sup> Regardless, however inaccessible Facebook may have made this data by putting it in cold storage or failing to cease the Hive Anon process, Facebook did so at its own

---

<sup>3</sup> The most generous explanation is that Mr. Duffey, who testified that he was prepared for his June 2 deposition in "~38 hours of meetings; 10 sessions; over 4 weeks," was not in touch with the lawyers who prepared the May 30 submission to the Special Master. Ex. 386 at -00003256.

<sup>4</sup> Mr. Duffey consulted with Ms. Ji in preparation for his deposition. Ex. 386 at -00003256. Ms. Allen is the most knowledgeable person about the Switchboard tool and is an E-Discovery team member. Duffey 30(b)(6) Dep. Rough Tr. at 29:19-20, 44:17-18. Plaintiffs propose that each of them be deposed on these topics immediately.

risk, and it should bear all costs of re-identifying and producing all Named Plaintiff data in the 137 Hive tables and any other tables that contain Named Plaintiff data, whether or not in cold storage. Plaintiffs thus propose that the process begin with a search of the 137 Hive tables, and that Facebook produce schemas, protocols and fields for those tables.<sup>5</sup> In addition, Plaintiffs propose that Facebook provide further explanation of the Hive namespaces in Exhibit A to the May 30 submission and the schema/fields of all Hive tables and internal developer documentation sufficient to explain what is in the fields.

**B. Facebook Must Immediately Produce All Named Plaintiff Data, Including Snapshots, and Certify That It Has Produced All Such Data That It Has Preserved.**

The preservation deposition revealed that Facebook has also misstated its ability to search for Named Plaintiffs' data. For the first time, Facebook admitted that the tool it regularly uses to preserve and produce information when it receives a subpoena about a user was not used to produce Named Plaintiffs' data in this case. Duffey 30(b)(6) Dep. Rough Tr. at 32:21-33:6. Facebook has repeatedly, and falsely, affirmed that the "DYI file for each individual user represents the most complete and best compilation of data Facebook maintains associated with that user." *See* Exhibit A. Plaintiffs request production of all withheld snapshots and a deposition of someone knowledgeable to explain why those files were withheld and what else exists.

**C. Facebook Must Produce Evidence Showing the Apps Installed by the Named Plaintiffs' Friends, or Else Formally Admit That It Cannot Do So.**

It is essential that the Named Plaintiffs learn which apps their friends installed—otherwise, it may not be possible to determine what third parties accessed the Named Plaintiffs' data, what data was accessed, and how it was used. For that reason, Facebook must produce evidence showing which apps were installed by the Named Plaintiffs' friends. If it cannot do so, it must admit—in a

---

<sup>5</sup> Facebook proposes the unworkable idea that responsive evidence should be selected only if it is "statistically significant." Proposing this metric is baffling, as Facebook has provided no information about the amount or substance of the data in each table. Counting the number of tables and calculating the number that equals ten percent of the total has zero probative value. This is a privacy class action in which the key questions are (1) *what* information was collected about users (the magnitude, subject matter and scope of collection); and (2) what Facebook did with it. With these underlying facts established, the parties will litigate whether that conduct was disclosed and whether Facebook obtained knowing consent from users.

binding and admissible manner—that it is incapable of doing so.<sup>6</sup>

### III. PLAINTIFFS’ PROPOSED RELIEF

First, by June 10, 2022, Facebook must produce all collected and/or preserved data relating to the Named Plaintiffs, including but not limited to all Switchboard snapshots, and identify, in writing and with specificity, any collections it seeks to withhold.

Second, by June 17, 2022, Facebook must produce all schemas, fields, and protos for the 137 Hive tables identified in the Duffey 30(b)(6) deposition and search these 137 tables using all Named Plaintiffs’ identifiers, including UIDs, RIDs, device IDS, and any other identifiers.

Third, by June 10, 2022, Facebook must provide information sufficient to describe the contents of the namespaces identified in Exhibit A to its May 30 submission, and provide schema, fields and protocols, or internal developer documentation to describe the contents of the Hive tables similar to Ex. 334 of the Clark 30(b)(6) Deposition.

Fourth, Mengge Ji and Jennifer Allen should sit for two 3-hour depositions next week.

Fifth, all Gibson Dunn partners engaged at any time on this matter, as well as in-house counsel Sandeep Solanki, Natalie Naugle, and Ian Chen, must submit affidavits attesting to whether Facebook has produced all ESI, data, content, and information for the Named Plaintiffs that it has preserved, including explanations as to why Facebook has failed to identify or produce the preserved 137 Hive tables or the snapshots to date.<sup>7</sup>

Plaintiffs may seek evidentiary preclusion orders based on Facebook’s misconduct.

---

<sup>6</sup> To be clear, other questions about Plaintiffs’ data remains, including where partner categories are stored; what happened to data for deprecated systems; and other issues.

<sup>7</sup> In addition to Gibson Dunn lawyers, Mr. Duffey identified Sandeep Solanki and Natalie Naugle as in-house counsel involved in preservation of ESI in response to the filing of this lawsuit. Duffey 30(b)(6) Dep. Rough Tr. 7:22-8:18.

Dated: June 6, 2022

Respectfully submitted,

KELLER ROHRBACK L.L.P.

BLEICHMAR FONTI & AULD LLP

By: /s/ Derek W. Loeser  
Derek W. Loeser

By: /s/ Lesley E. Weaver  
Lesley E. Weaver

Derek W. Loeser (admitted *pro hac vice*)  
Cari Campen Laufenberg (admitted *pro hac vice*)  
David Ko (admitted *pro hac vice*)  
Adele A. Daniel (admitted *pro hac vice*)  
Benjamin Gould (SBN 250630)  
1201 Third Avenue, Suite 3200  
Seattle, WA 98101  
Tel.: (206) 623-1900  
Fax: (206) 623-3384  
dloeser@kellerrohrback.com  
claufenberg@kellerrohrback.com  
dko@kellerrohrback.com  
adaniel@kellerrohrback.com  
bgould@kellerrohrback.com

Lesley E. Weaver (SBN 191305)  
Anne K. Davis (SBN 267909)  
Matthew S. Melamed (SBN 260272)  
Angelica M. Ornelas (SBN 285929)  
Joshua D. Samra (SBN 313050)  
555 12th Street, Suite 1600  
Oakland, CA 94607  
Tel.: (415) 445-4003  
Fax: (415) 445-4020  
lweaver@bfalaw.com  
adavis@bfalaw.com  
mmelamed@bfalaw.com  
aornelas@bfalaw.com  
jsamra@bfalaw.com

Christopher Springer (SBN 291180)  
801 Garden Street, Suite 301  
Santa Barbara, CA 93101  
Tel.: (805) 456-1496  
Fax: (805) 456-1497  
cspringer@kellerrohrback.com

Eric Fierro (admitted *pro hac vice*)  
3101 North Central Avenue, Suite 1400  
Phoenix, AZ 85012  
Tel: (602) 248-0088  
Fax: (602) 248-2822  
efierro@kellerrohrback.com

*Plaintiffs' Co-Lead Counsel*

**ATTESTATION PURSUANT TO CIVIL LOCAL RULE 5-1(h)(3)**

I, Lesley E. Weaver, attest that concurrence in the filing of this document has been obtained from the other signatory. I declare under penalty of perjury that the foregoing is true and correct.

Executed this 6th day of June, 2022, at Oakland, California.

/s/ Lesley E. Weaver  
Lesley E. Weaver

# EXHIBIT A

**EXHIBIT A**  
**CHART OF FACEBOOK’S MISSTATEMENTS TO THE SPECIAL MASTER AND TO THE COURT  
REGARDING THE NAMED PLAINTIFFS’ DATA**

<b>Date</b>	<b>Court Submission</b>	<b>Facebook’s (False) Representations About the Named Plaintiffs’ Data</b>	<b>Testimony of Facebook’s 30(b)(6) Deponents<sup>1</sup></b>
August 13, 2020	Joint Status Update	Facebook states: “To find the data Plaintiffs seek, Facebook would need to identify every single internal analysis that uses Platform data and attempt to de-anonymize every data point within those analyses to determine if any provided by the Named Plaintiffs is among the data. Even a large team of engineers working full time for several years likely could not identify all of the information Plaintiffs seek.” Dkt. No. 495 at 6-7.	52 snapshots taken with the Switchboard tool are in Facebook’s possession but have not been produced. 137 Hive tables relating to the litigation were set aside and preserved. <i>See</i> Duffey 30(b)(6) Dep. Rough Tr. at 153:14-16 (“[W]e have 137 Hive data tables preserved and on legal hold for this matter.”); <i>Id.</i> at 147:16-19 (“We’ve taken snapshots indirectly the Switchboard tool and the DYI tool [] for each of the named plaintiffs that had data available at the time of the snapshot.”); Ex. 386 at ADVANCE-META-00003259 (“137 Hive tables on legal hold.”).  Neither the snapshots nor the preserved 137 Hive tables were disclosed until last Thursday, June 2, 2022 in a 30(b)(6) deposition. It would not have taken a large team of engineers working full time for several years to search and produce Named Plaintiffs’ data contained in them.
October 8, 2020	Facebook’s Reply Brief ISO Request to Enforce the Partial Stay of Discovery in PTO No. 20	Facebook argued that Plaintiffs’ request for additional data relating to the Named Plaintiffs’ concerned data outside the scope of any live theory of this case and would require it to search “millions of disaggregated data sets.” Dkt. No.	52 snapshots taken with the Switchboard tool are in Facebook’s possession but have not been produced. <i>See</i> Duffey 30(b)(6) Dep. Rough Tr. at 147:16-19 (“We’ve taken snapshots indirectly the

<sup>1</sup> Plaintiffs will provide the Special Master with transcripts upon request.

Date	Court Submission	Facebook's (False) Representations About the Named Plaintiffs' Data	Testimony of Facebook's 30(b)(6) Deponents <sup>1</sup>
		537 at 6. It also stated that additional data "cannot reasonably" be identified. <i>Id.</i> at 9.	<p>Switchboard tool and the DYI tool [] for each of the named plaintiffs that had data available at the time of the snapshot."); Ex. 386 at ADVANCE-META-00003258 ("In the aggregate, FB preserved 52 Switchboard snapshots across Groups, Pages, Ads Accounts in Fall 2021 for the named plaintiffs when identified as Admins.").</p> <p>137 Hive tables relating to the litigation were set aside and preserved. <i>See</i> Duffey 30(b)(6) Dep. Rough Tr. at 153:14-16 ("[W]e have 137 Hive data tables preserved and on legal hold for this matter."); Ex. 386 at ADVANCE-META-00003259 ("137 Hive tables on legal hold."). Facebook did not disclose this to the Special Master, the Court, or Plaintiffs in any of the multiple hearings and briefings on this issue, and it was not revealed until last Thursday, June 2, 2022 in a 30(b)(6) deposition.</p>
October 8, 2020	Facebook's Reply Brief ISO Request to Enforce the Partial Stay of Discovery in PTO No. 20	"There is no way for Facebook to run a centralized search for a user's ID, random ID, or any 'hashed data' identifiers across millions of data sets[.]" Dkt. No. 537 at 10.	<p>"[I]f I know the table and in working with a specific table, yes, I can query for a specific RID or user ID." Clark 30(b)(6) Dep. Tr. at 99:9-18.</p> <p>By failing to disclose that Facebook had already identified 137 Hive tables relevant</p>



Date	Court Submission	Facebook's (False) Representations About the Named Plaintiffs' Data	Testimony of Facebook's 30(b)(6) Deponents <sup>1</sup>
			to the litigation, which can be searched for UID and RID, among other identifiers, Facebook delayed and complicated the search for and production of Named Plaintiffs' data.
October 18, 2021	Decl. of Mengge Ji ISO Facebook Inc.'s Motion for a Protective Order Against Production of API Call Logs	"Tables in cold storage cannot be reviewed or analyzed while they remain in cold storage." Ji Decl. at ¶ 18.	If the Hive tables were put in cold storage for purposes of preserving them for litigation, then Facebook must be able to produce data from those tables. <i>See</i> Duffey 30(b)(6) Dep. Rough Tr. at 175:10-13 ("I am aware of our E-Discovery data science team producing structured data ... in matters. Whether or not it was taken out of cold storage or not, I'm ... not clear on that process."); <i>Id.</i> at 176:7-13 ("When data is put into cold storage, only the E-Discovery team would ... have access to that data.").
October 28, 2021	Facebook's Separate Statement in Support of Facebook's Opposition to Plaintiffs' Motion to Compel Production of Named Plaintiffs' Content and Information	"Facebook subsequently repeatedly informed Plaintiffs and the Court that it had produced the most complete compilation of data associated with the Named Plaintiffs' accounts, including on July 30, 2020 (Dkt. 484), August 13, 2020 (Dkt. 495), and September 18, 2020 (Dkt. 515)." Facebook's description of its own statements to the Court, Separate Stmt. at 2.	These repeated statements were false. The DYI files do not contain information about users' privacy settings, Pages, Groups, or adversities interests that the Switchboard tool does, including about the "relationships between a user" and information about who the user is interacting with. The files are also in a less usable format than the Switchboard

Date	Court Submission	Facebook's (False) Representations About the Named Plaintiffs' Data	Testimony of Facebook's 30(b)(6) Deponents <sup>1</sup>
October 28, 2021	Declaration of Ben Mitchell ISO Facebook's Opp. to Pls' MTC Production of Named Plaintiffs' Content and Information	"The DYI file for each individual user represents the most complete and best compilation of data Facebook maintains associated with that user, and the best available compilation of the data about that user in the Social Graph, in a human-readable and producible form." Mitchell Decl. at ¶ 5.	<p>snapshots. Duffey 30(b)(6) Dep. Rough Tr. at 56:5-8; <i>see also</i>:</p> <ul style="list-style-type: none"> <li>• "Switchboard has some additional information that DYI doesn't pertaining to Facebook groups and advertising accounts." <i>Id.</i> at 31:11-13;</li> <li>• "I know that Switchboard was a tool used by that team before DYI. ... [T]he law enforcement response team can produce records to law enforcement in response to subpoena in a more useable way than DYI has available to it as a download. ... In a, you know, like -- a PDF record." <i>Id.</i> at 32:21-33:9;</li> <li>• "I believe that Switchboard captures additional information about who the user is interacting with." <i>Id.</i> at 45:7-10;</li> <li>• "I do believe that the Switchboard does collect ... additional information potentially from TAO that is not collected in a DYI snapshot." <i>Id.</i> at 56:9-12;</li> <li>• "I believe there are privacy settings captured in Switchboard[.]" <i>Id.</i> at 143:22-23;</li> <li>• "Switchboard: Has more information about other users</li> </ul>
November 15, 2021	Facebook's Response to Objection Regarding Named Plaintiffs' Data Briefing	<p>"Facebook explained repeatedly that the Download Your Information ('DYI') files it produced satisfied its production obligations under Discovery Order 9 and the technical reasons why this is so." Facebook's Response at 1.</p> <p>"And a user's DYI file contains a human-readable download of the most complete set of data about that user in the Social Graph (and more)." <i>Id.</i> at 2.</p>	

Date	Court Submission	Facebook's (False) Representations About the Named Plaintiffs' Data	Testimony of Facebook's 30(b)(6) Deponents <sup>1</sup>
			(compared to DYI). It also has information related to Pages, Groups, Advertising accounts." Ex. 386 at ADVANCE-META-00003258.
December 10, 2021	Decl. of David Pope ISO Facebook's Mot. for Reconsideration of the Special Master's Order Re: Named Plaintiffs' MTC Production of Plaintiff Data	In response to the Special Master's request for Facebook to provide "'a list of data sources that may contain' user data," Facebook submitted the declaration of David Pope identifying "149 data systems ... identified as containing user data[.]" Pope Decl. at ¶¶ 5, 9.	This submission was grossly misleading, including because it failed to identify the 52 snapshots taken with the Switchboard tool of the Named Plaintiffs or the 137 Hive tables relating to the litigation were set aside and preserved. Duffey 30(b)(6) Dep. Rough Tr. at 153:14-16 ("[W]e have 137 Hive data tables preserved and on legal hold for this matter."); <i>Id.</i> at 147:16-19 ("We've taken snapshots indirectly the Switchboard tool and the DYI tool [] for each of the named plaintiffs that had data available at the time of the snapshot.").
December 10, 2021	Decl. of Mengge Ji ISO Facebook's Mot. for Reconsideration of the Special Master's Order Re: Named Plaintiffs' MTC Production of Plaintiff Data	"Facebook cannot quickly or efficiently run searches, for instance, for particular user IDs, across multiple days of data in a partitioned table at once." Ji Decl. at ¶ 21.	<p>"I can search a table, if I know a specific table and I'm looking for a specific RID[.]" Clark 30(b)(6) Dep. Tr. at 87:19-21.</p> <p>"[I]f I know the table and in working with a specific table, yes, I can query for a specific RID or user ID." <i>Id.</i> at 99:9-18.</p> <p>Facebook did not tell the Special Master that it could identify specific tables and</p>

Date	Court Submission	Facebook's (False) Representations About the Named Plaintiffs' Data	Testimony of Facebook's 30(b)(6) Deponents <sup>1</sup>
			that it could search using Named Plaintiffs' RIDs and UIDs, among other identifiers.
December 22, 2021	Motion for Leave to File and For Stay Pending Appeal of the Special Master's Amended Order Re: Plaintiffs' Motion to Compel Production of Plaintiff Data	In response to the Special Master's Initial Order requiring Facebook to provide "a list of data sources related to the Named Plaintiffs," Facebook provided a list of "potential data sources but explained (and substantiated with uncontroverted declarations) that compiling the remaining information would take more than one year of work and would require coordination across dozens of Facebook teams and hundreds of Facebook employees." Mot. for Leave at 3-4.	These statements were misleading. Facebook did not inform the Special Master about either the 137 Hive tables preserved for this litigation, or the 52 snapshots preserved using the Switchboard tool. <i>See</i> Duffey 30(b)(6) Dep. Rough Tr. at 60:21-25 ("I understand that there are 137 Hive tables placed on legal hold for the Cambridge Analytica matter."); <i>Id.</i> at 50:16-20 ("If any of the named plaintiffs were admins for any groups or pages, then snapshots were taken of those groups and pages, and if any of the named plaintiffs had advertising accounts, those were also captured in a Switchboard snapshot.").
March 7, 2022	Letter from Rose Ring in Response to Special Master's February 21 Order re: Plaintiff Data	"[T]he DYI system includes all three categories of data: (1) data collected from a user's on-platform activity, (2) data obtained from third parties regarding a user's off-platform activities, and (3) data inferred from a user's on or off-platform activity. So the issue is whether Facebook should produce more data in categories (1), (2), and (3) from systems other than DYI, which we understand is the focus of the Special Master's remaining questions." Mar. 7, 2022 Ltr. from R. Ring at 1.	The reason the Special Master was focusing on the DYI tool is because Facebook had repeatedly told the Court, incorrectly, that it was a complete source for user information.
March 9, 2022	Special Master Hearing	"And so there becomes a technical feasibility because not all that storage that sits in Hive is online. And so that would – it becomes a technical feasibility challenge to do that full discovery	"Scuba is a tool to help data scientists and data engineers, where they know where data lives in Hive, to be able to write a query to do analysis or -- or analytics on a

Date	Court Submission	Facebook's (False) Representations About the Named Plaintiffs' Data	Testimony of Facebook's 30(b)(6) Deponents <sup>1</sup>
		across everything." Mar. 9, 2022 Special Master Hearing Tr. at 115:2-9.	subset of data." Clark 30(b)(6) Dep. Tr. at 89:19-23.
April 11, 2022	Letter from Rose Ring in Response to Special Master's March 9 Order re: Plaintiff Data	Facebook responded to the Special Master's request for information about Hive tables, and identified "Hive tables that store user identifiers with data as it enters Hive." Apr. 11, 2022 Ltr. from R. Ring at 2.	<p>"I understand that there are 137 Hive tables placed on legal hold for the Cambridge Analytica matter." Duffey 30(b)(6) Dep. Rough Tr. at 60:21-25.</p> <p>"Hive data that is on legal hold to prevent it from being deleted or modified." <i>Id.</i> at 64:2-3.</p> <p>"We're talking about 137 Hive tables. I would imagine that Hive tables were put on legal hold starting in 2018 and have, you know, as additional -- you know, again, as additional Hive tables are identified relevant to a matter, we work with our -- our in-house and outside counsel and our data science team and E-discovery to place those Hive tables on legal hold." <i>Id.</i> at 66:18-67:1.</p>

# Exhibit W

June 16, 2022

VIA JAMS ACCESS

Client: 30993-00116

Dear Special Master Garrie,

Below Facebook responds to the Special Master's JAMS message dated June 10, 2022 stating "Facebook is to respond to the issues regarding production of Named Plaintiff data in Switchboard and the 137 Hive tables raised in Plaintiff's June 6, 2022 brief."

After months of proceedings, the parties submitted proposals regarding the production of additional Named Plaintiff data. The Special Master held a hearing on May 17, 2022, during which he identified three areas of disagreement. On May 24, 2022, the Special Master ordered the parties to submit simultaneous briefing on those issues, but rather than address them, Plaintiffs focused on new issues based on supposed "new evidence." Plaintiffs' "new evidence" is largely irrelevant to these proceedings, and Plaintiffs also misconstrue and misinterpret the testimony they cite. Specifically, Plaintiffs point to data a witness testified was preserved from a system called Switchboard and Hive. Even though the witness did not testify that the preserved data is Named Plaintiff data or Named Plaintiff data that has not been made available to Plaintiffs, they assume that is the case and based on those assumptions make false accusations and an unprecedented demand for affidavits from more than a dozen attorneys. Facebook addresses each of these issues in turn below.

### 1. Switchboard

Switchboard is a tool created to respond to law-enforcement subpoenas, <https://www.facebook.com/safety/groups/law/guidelines/>, which are not at issue in this case. Law enforcement can request information in response to lawful subpoenas, through Facebook's Law Enforcement Online Request System. Law enforcement subpoenas must be narrowly tailored and particularized, and the relevant categories of data (and only those) are provided through Switchboard. Switchboard is not designed or intended for use by Facebook users or in private litigation. That said, Facebook preserved data from Switchboard in connection with this case out of an abundance of caution.

To be clear, Switchboard data largely overlaps with data already produced. Nonetheless, Facebook agrees to produce Named Plaintiff data preserved from Switchboard for current named plaintiffs. Facebook addresses Plaintiffs' accusations below:

**DYI.** Plaintiffs say testimony about Switchboard undermines statements that DYI is the "most complete compilation of data associated with the Named Plaintiffs' accounts" because there is information in Switchboard that is not in DYI. There is no inconsistency.

DYI is the "most complete compilation of data associated with the Named Plaintiffs' accounts." In fact, DYI contains many types of data that are not in Switchboard. To name just a few: "your topics," which is a collection of topics determined by a user's activity that is used to create recommendations for users in different areas of Facebook; friend peer group; creator badges (including labels like "visual storyteller" or "conversation starter" based on activity in Groups); inferred language information; users a user has chosen to "see less" or "see first" in News Feed; and music recommendations based a user's interactions.

Data associated with a user's account preserved in a Switchboard snapshot largely tracks a subset of the data in DYI or otherwise available to the user. In some instances data in Switchboard and DYI appear in a different format or with different details, given their different uses. Unlike DYI, a Switchboard snapshot contains (1) data about other users and

(2) data created for law enforcement (neither of which are relevant to these Named Plaintiff data proceedings). For (1) an example is that a Switchboard snapshot includes information about users who reported or blocked the account and information about other users' posts and IP addresses. For (2) an example is that Facebook analyzes illegal child exploitive content and creates reports for the National Center for Missing and Exploited Children.

In Discovery Order No. 9, Judge Corley identified three categories of "discoverable user data" about the Named Plaintiffs: (1) data collected from a user's on-platform activity, (2) data obtained from third parties regarding a user's off-platform activities, and (3) data inferred from a user's on or off-platform activity. Data about other users and analyses prepared for law enforcement do not fall into these categories. Facebook also cannot produce other users' data without their consent, consistent with its obligations to those users, including under the Stored Communications Act.

**Privacy Settings.** Plaintiffs say Facebook withheld privacy settings from preserved Switchboard snapshots. This is false. Facebook produced snapshots of the Named Plaintiffs' privacy settings from Switchboard (attached as **Exhibit A**).

**Other Users.** Plaintiffs say that, unlike DYI, Switchboard contains information about "the relationships between a user." As explained above, the Switchboard snapshot of an account contains information about other users that is not Named Plaintiff data. This data does not fall into any of the three categories identified in Discovery Order No. 9.

**Groups and Pages.** Plaintiffs say Facebook withheld Named Plaintiff data from the Switchboard snapshots about Groups and Pages. This is wrong. A user's DYI file contains information about their own activity in Groups and on Pages. The Switchboard snapshots contain additional information about *other users'* activities in Groups and on Pages.

**Advertising Information.** Plaintiffs suggest Facebook withheld Named Plaintiff data from Switchboard snapshots about advertisements the Named Plaintiffs received. This is also wrong. Like DYI, a Switchboard snapshot of a user's account contains ad interests and ads with which a user interacted. The cited testimony refers to snapshots of advertising accounts, which are part of a Business Page, and largely contains data relating to ads placed by that Page. This is a consumer class action, and Business Pages are not at issue. In any case, the categories of data in a snapshot of an advertising account overlap almost entirely with categories of data in DYI, such as the Page administrators' (of which there are often more than one) Names, IP Addresses, Phone Numbers, Email Addresses, and Logins. The additional data is payment/tax information in connection with purchasing advertisements and the actual ad campaigns run; that information is available to an advertiser.

**Format.** Plaintiffs say deposition testimony shows a Switchboard snapshot is "more usable" than DYI. Plaintiffs misconstrue the testimony. The witness explained the specific categories of Switchboard data provided in response to a subpoena can be produced to law enforcement as a .pdf, while DYI is "lots of individual files." Trans. at 51:19-23. DYI contains multiple files so a user can download individual categories and easily navigate the results. For litigation, individual files are also more usable. Most DYI files produced in this case are tens of thousands of pages; some are hundreds of thousands of pages. If Facebook produced the data as a single .pdf, it would be extremely difficult to navigate and use.

\* \* \*

Again, Facebook will produce the Named Plaintiff data preserved in Switchboard for current Named Plaintiffs. As disclosed previously, Facebook added new categories of data to the DYI system since producing certain Named Plaintiffs' DYI data. Facebook explained months ago that it is "happy to produce updated versions of DYI files . . . but has understood that Plaintiffs are not interested in additional DYI data, since it is 'already available to Plaintiffs.'" See, e.g. Dkt. 911 at 35 n.22. Although Plaintiffs have not responded, Facebook will produce updated DYI files, to avoid any confusion as to what data is available to users.



## 2. Hive Tables

Facebook's corporate designee testified that Facebook preserved various Hive tables in connection with this and related litigation. Plaintiffs leap to the conclusion that all of these tables contain Named Plaintiff data and say Facebook improperly failed to disclose the tables. Again, even though Plaintiffs' assumptions are wrong, Facebook agrees to produce the schema (i.e. column names) for preserved tables containing user identifiers and will meet and confer with Plaintiffs regarding a production of Named Plaintiff data from those tables. This overlaps with information and data Facebook already agreed to provide. Below Facebook addresses Plaintiffs' accusations about preserved Hive tables.

**Named Plaintiff Data.** Plaintiffs say Facebook should have disclosed a list of preserved Hive tables because such a list would have identified the most relevant Named Plaintiff data. This is wrong. Tables in Hive were preserved in connection with this litigation and related litigation for a number of reasons, many of which have nothing to do with Named Plaintiff data. Of the 137 tables identified, a minority contain user identifiers, and several of those tables were put on hold in connection with Facebook's April 18, 2022 proposal to produce categories of data requested or referenced by Plaintiffs.

**Searchability.** Plaintiffs say testimony that Hive tables can be searched by UID or RID undermines prior representations that Plaintiffs' requests would require searching "millions of disaggregated data sets." There is no inconsistency.

The cited statements concern a broader demand that Plaintiffs narrowed. Plaintiffs initially requested every data point at Facebook that might relate back to a Named Plaintiff in any way, including aggregated/fully anonymized data. Facebook explained that satisfying this request would be technically infeasible as it would require reverse engineering the underlying sources of data in millions of tables that are not reasonably capable of being identified to a user. For example, a table could log the number of Facebook users who log into Facebook each hour, to track peak activity for capacity budgeting, without any data about which users logged in each hour. Given the breadth of the request, Facebook explained "[e]ven a large team of engineers working full time for several years likely could not" "identify all of the information plaintiffs seek." That is true.

Before the Special Master, the parties have discussed a narrower set of data that is capable of being associated with a user's account—tables that contain user identifiers.

**Cold Storage.** Plaintiffs say testimony that "offline" tables can be searched undermines representations about the searchability of data in cold storage. The testimony referred to Hive, not cold storage. As Facebook has explained, Hive is "a data system that is used for offline data storage and analysis," *see* 5/11/2022 Submission p. 2, unlike TAO, which supports the online production environment. Plaintiffs also say testimony undermines representations that data in cold storage cannot be produced. Facebook did not represent that data in cold storage cannot be produced. It argued it should not be required to restore tables consisting of hundreds of petabytes of data, in part because it had accessible (and agreed to produce), summary versions of the data. The Special Master agreed. *See* Nov. 8, 2021 Order Re: Facebook's Motion For Protective Order Against Production of API Call Logs.

## 3. Request for Affidavits from Counsel

Based on the accusations above, Plaintiffs make an unprecedented demand for affidavits from Facebook's in-house legal team and Gibson Dunn partners to have worked on this matter regarding Facebook's production of preserved Named Plaintiff data. There is no basis for a request of this nature, which is based on Plaintiffs' mischaracterizations of the record, and has nothing to do with the outstanding areas of dispute regarding the parties' Named Plaintiff data proposals. In any case, Plaintiffs received corporate testimony about Facebook's preservation efforts, and, as explained above, agrees to produce data preserved in sources to which the witness testified. This request should be denied.

Sincerely,

/s/ Heather L. Richardson

Heather L. Richardson

# Exhibit X

Derek W. Loeser (admitted *pro hac vice*)  
KELLER ROHRBACK L.L.P.  
1201 Third Avenue, Suite 3200  
Seattle, WA 98101  
Tel.: (206) 623-1900  
Fax: (206) 623-3384  
dloeser@kellerrohrback.com

Lesley E. Weaver (SBN 191305)  
BLEICHMAR FONTI & AULD LLP  
555 12th Street, Suite 1600  
Oakland, CA 94607  
Tel.: (415) 445-4003  
Fax: (415) 445-4020  
lweaver@bfalaw.com

*Plaintiffs' Co-Lead Counsel*

*Additional counsel listed on signature page*

**UNITED STATES DISTRICT COURT  
NORTHERN DISTRICT OF CALIFORNIA  
SAN FRANCISCO DIVISION**

IN RE: FACEBOOK, INC. CONSUMER  
PRIVACY USER PROFILE LITIGATION

MDL No. 2843  
Case No. 18-md-02843-VC

This document relates to:

ALL ACTIONS

**PLAINTIFFS' FURTHER RESPONSE  
RE: NAMED PLAINTIFF DATA  
ADDRESSING RECENTLY REVEALED  
PRESERVED DATA**

Judge: Hon. Vince Chhabra  
Special Master Daniel Garrie  
Courtroom: 4, 17th Floor

JAMS Ref. No.: 1200058674

PLS' FURTHER RESPONSE RE: NAMED  
PLAINTIFF DATA ADDRESSING  
RECENTLY REVEALED PRESERVED  
DATA

MDL No. 2843  
CASE No. 18-MD-02843-VC  
JAMS REF. No.: 1200058674

This brief is submitted pursuant to the Special Master's June 10 and 14 emails which granted Facebook the right to respond to Plaintiffs' inquiries regarding Named Plaintiff data in Switchboard and 137 Hive tables, and Plaintiffs the right to submit a brief by June 20, 2022.

Facebook's letter of June 16, 2022 offers argument from counsel about what is contained in Switchboard snapshots with few citations to evidence, testimony, or documents. In contrast, Plaintiffs' submission rested entirely on the sworn testimony of Facebook's corporate representative, Michael Duffey, who proffered binding testimony on behalf of Facebook regarding what was preserved by Facebook in this action after March 2018 when the Cambridge Analytica scandal broke. Mr. Duffey testified:

- "We've taken snapshots in the Switchboard tool and the DYI tool [] for each of the named plaintiffs that had data available at the time of the snapshot." Duffey 30(b)(6) Dep. Tr. at 158:1-4;
- "Switchboard has some additional information that DYI doesn't pertaining to Facebook groups and advertising accounts." *Id.* at 40:4-6;
- "I believe that Switchboard captures additional information about who the user is interacting with." *Id.* at 53:25-55:2;
- "I do believe that the Switchboard does collect ... additional information potentially from TAO that is not collected in a DYI snapshot." *Id.* at 65:1-4;
- "I believe there are privacy settings captured in Switchboard[.]" *Id.* at 154:5-6.

Moreover, ahead of his testimony, Mr. Duffey prepared notes which describe the Switchboard snapshots as "ha[ving] more information about other users (compared to DYI)" and "also ha[ving] information related to Pages, Groups, Advertising accounts." Ex. 386 at ADVANCE-META-00003258. Thus, when Facebook's letter claims throughout that "Plaintiffs suggest" and "Plaintiffs say," this is erroneous. Plaintiffs did not say; Facebook did, as shown by the sworn testimony of Facebook's corporate representative. *Compare* June 16, 2022 Letter from H. Richardson re Named Pls' Data at 2 ("June 16 Letter") ("Plaintiffs say deposition testimony shows a Switchboard snapshot is 'more usable' than DYI") *with* Duffey 30(b)(6) Dep. Tr. at 41:20-24 ("[T]he law enforcement response team can produce records" from the Switchboard tool "to law enforcement in response to subpoena in a more useable way than DYI has available to it as a download.").

Facebook offers no explanation for its failure to identify these files, let alone produce them, years ago. Facebook claims that "the relationships between a user" is not "Named Plaintiff data." June 16 Letter at 2 (with no citation). But clearly actions between the Named Plaintiffs and other users would encompass friends, and the unconsented-to sharing of friend data is centrally relevant to this case. Facebook appears to be saying that none of the Switchboard snapshots' information about relationships between Named Plaintiffs and users reflects (1) data collected from a user's on-platform activity, (2) data obtained from third parties regarding a user's off-platform activities, or (3) data inferred from a user's on or off-platform activity. This seems unlikely. In any event,

Facebook has agreed to produce “the Named Plaintiff data preserved in Switchboard for current Named Plaintiffs.” *Id.* at 2. It is unclear if Facebook is referring to all 52 snapshots preserved over time for the Named Plaintiffs, or if that is in addition to other data preserved in Switchboard. Plaintiffs seek both and request that the Special Master order production by Friday, June 24, 2022.<sup>1</sup>

With regard to the Hive tables, Facebook is equally opaque, citing no documents or evidence. Mr. Duffey testified that Facebook has “137 Hive data tables preserved and on legal hold for this matter.” Duffey 30(b)(6) Dep. Tr. at 163:24-164:1; *see also* Ex. 386 at ADVANCE-META-00003259 (confirming “137 Hive tables on legal hold.”). Facebook offers no explanation for why it did not earlier disclose the 137 Hive tables, while saying it might produce column names for preserved tables containing user identifiers. Facebook offers no date by which it will do so. Facebook writes that Plaintiffs “assume” that the Tables relate to the Named Plaintiffs. This is incorrect. The parties have been working with the Special Master for months to find an efficient way to identify Hive tables to search for Named Plaintiff data. And FB-CA-MDL-03452017 (attached as Exhibit A) shows that Facebook creates descriptions of Hive tables in the regular course of business. Pages 4-5 of this document provide an example of what Plaintiffs seek here.<sup>2</sup> At no point prior to the deposition of Mr. Duffey did anyone from Facebook, including counsel, inform the Special Master or Plaintiffs that Facebook had *already identified* 137 Hive tables and preserved them for this litigation. They are the obvious place to start searching.

Plaintiffs request an order requiring a production date of Friday, June 24, 2022 for the Hive schema, fields and documents sufficient to describe the contents of the tables. Plaintiffs further request an order to show cause why Facebook should not be ordered to produce all Hive tables referencing user identifiers immediately, given that Facebook did not earlier advise the Special Master or Plaintiffs of their existence.

Because the representations of Facebook’s counsel in their briefing are inconsistent with Facebook’s sworn testimony, Plaintiffs propose that affidavits are an appropriate remedy to ensure that Facebook has accurately identified all potential sources of Named Plaintiff data.

### **Cold Storage**

With regard to searchability and cold storage, Facebook’s letter continues to aver that “Plaintiffs say,” and the letter ignores Facebook’s own sworn testimony: “I can search a table, if I know a specific table and I’m looking for a specific RID[.]” Mike Clark 30(b)(6) Dep. Tr. at 87:19-21; *see also* Ex. 386 at ADVANCE-META-00003259 (“iData: Allows searching for Hive tables and indicates if table is on hold.”). Facebook’s letter also ignores that it placed the 137 Hive tables in cold storage only after identifying them as relevant to this case. *See* Ex. 386 at ADVANCE-META-00003259.

---

<sup>1</sup> Plaintiffs reserve the right to respond once the Snapshots and DYI files are produced and to seek further corporate testimony from Facebook once the documents are produced.

<sup>2</sup> This document was improperly withheld by Facebook as privileged, and then de-designated and produced on June 6, 2022, the same day as Plaintiffs’ June 6 submission.

Facebook refers to the Special Master's ruling regarding certain tables in cold storage, but that ruling related to different content, and the parties and the Special Master have learned more now. FB-CA-MDL-03330509 (attached as Exhibit B) raises questions regarding Facebook's representations throughout the inquiry into Named Plaintiff data concerning the burden of restoring information in cold storage.

The document is an internal chat among a large number of Facebook employees, including custodians Eugene Zarakhovsky (now Zarashaw) and Shirine Sajjadi. Towards the beginning of the chat, an employee states: "Data Infra now says they can restore all the relevant partitions of [REDACTED] in \*72 hours\* if we throw everything at it, which is very different from 150 days quoted earlier." Ex. B at FB-CA-MDL-03330511. In other words, committing resources to restoring a table from cold storage would enable Facebook to restore the table 50 times quicker than initially estimated. Approximately 11 hours later, the same engineer updates the effort: "[T]he positively surprising news from the army of data infra working on cold storage is that we're very close to having full restore on the [REDACTED]. Way faster than 150 days, or 72 hours :) still working to make sure that is indeed correct, but it's looking like it." Ex. B at FB-CA-MDL-03330519. The restoration from cold storage, initially estimated to take 150 days, had been cut to approximately a half day—300 times faster than the initial estimate.

Similar to the previous example, Facebook improperly withheld this document as privileged, then de-designated and produced it only two weeks before Plaintiffs' June 6 submission. This document reveals that Facebook can restore data from cold storage in fairly short order. It should be ordered to do so.

Dated: June 20, 2022

Respectfully submitted,

KELLER ROHRBACK L.L.P.

BLEICHMAR FONTI & AULD LLP

By: /s/ Derek W. Loeser  
Derek W. Loeser

By: /s/ Lesley E. Weaver  
Lesley E. Weaver

Derek W. Loeser (admitted *pro hac vice*)  
Cari Campen Laufenberg (admitted *pro hac vice*)  
David Ko (admitted *pro hac vice*)  
Adele A. Daniel (admitted *pro hac vice*)  
Benjamin Gould (SBN 250630)  
1201 Third Avenue, Suite 3200  
Seattle, WA 98101  
Tel.: (206) 623-1900  
Fax: (206) 623-3384  
dloeser@kellerrohrback.com  
claufenberg@kellerrohrback.com  
dko@kellerrohrback.com

Lesley E. Weaver (SBN 191305)  
Anne K. Davis (SBN 267909)  
Matthew S. Melamed (SBN 260272)  
Angelica M. Ornelas (SBN 285929)  
Joshua D. Samra (SBN 313050)  
555 12th Street, Suite 1600  
Oakland, CA 94607  
Tel.: (415) 445-4003  
Fax: (415) 445-4020  
lweaver@bfalaw.com  
adavis@bfalaw.com  
mmelamed@bfalaw.com

adaniel@kellerrohrback.com  
bgould@kellerrohrback.com

aornelas@bfalaw.com  
jsamra@bfalaw.com

Christopher Springer (SBN 291180)  
801 Garden Street, Suite 301  
Santa Barbara, CA 93101  
Tel.: (805) 456-1496  
Fax: (805) 456-1497  
cspringer@kellerrohrback.com

Eric Fierro (admitted *pro hac vice*)  
3101 North Central Avenue, Suite 1400  
Phoenix, AZ 85012  
Tel: (602) 248-0088  
Fax: (602) 248-2822  
efierro@kellerrohrback.com

*Plaintiffs' Co-Lead Counsel*



**ATTESTATION PURSUANT TO CIVIL LOCAL RULE 5-1(h)(3)**

I, Lesley E. Weaver, attest that concurrence in the filing of this document has been obtained from the other signatory. I declare under penalty of perjury that the foregoing is true and correct.

Executed this 20th day of June, 2022, at Oakland, California.

/s/ Lesley E. Weaver  
Lesley E. Weaver